RJP

The Richmond Journal of **Philosophy**

Volume **One** Issue **Five** **Autumn 2003**

In this *issue...*

David Papineau
on consciousness

Will Cartwright
on freedom

John Collins
on language

Paul Sperring
on zombies

Kathy Behrendt
on personal identity

Christopher Norris
on anti-realism

Richmond upon Thames College

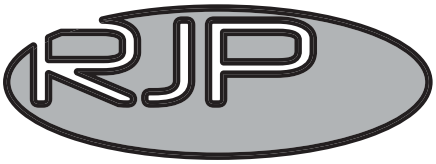

The Richmond Journal of Philosophy

Issue five
Autumn 2003

Philosophy Department

Richmond upon Thames College

Egerton Road

Twickenham

Middlesex

TW2 7SJ

United Kingdom

email: rjp@rutc.ac.uk

www.rutc.ac.uk/rjp

Contents

[Editorial]

Welcome to the fifth edition of the Richmond Journal of Philosophy. David Papineau considers some confusions about consciousness in the first article. In particular he addresses the view that our intuitions about mind-body dualism lead us astray in thinking that there must be something metaphysically 'special' about consciousness. For the second paper we move from a topic that has gained considerable attention in the philosophy of mind to one that continues to excite controversy in the domain of political philosophy. Will Cartwright assesses Mill's approach to freedom of discussion. Next John Collins provides in dialogue form an overview of Noam Chomsky's theory of language acquisition. In terms of its influence in linguistics and the philosophy of language it is hard to overestimate the significance of Chomsky's work. In the fourth of our papers Paul Sperring asks whether zombies are conceivable. He has in mind not the undead, but a person's physical replica which lacks any of the person's mental states. If zombies are conceivable then some will argue that dualism of some kind in the philosophy of mind can be defended. After zombies Kathy Behrendt turns our attention to the question of personal identity by examining the approaches of David Hume and the contemporary philosopher Derek Parfit. The paper does not simply set out their views but considers the tension that arises between philosophical theory and our everyday attitudes. In our final paper Chris Norris takes us to the heart of a debate that has characterised much of contemporary philosophy – the issues of anti-realism, scepticism and meaning.

Purpose of the Journal

The motivation for and ambition of the journal is to provide serious philosophy for students who are at an early stage in their philosophical studies. The style and content of the papers will be accessible to students who have yet to become hardened to the more technical and specialised journals of professional philosophy.

What do we mean by 'serious' philosophy? First, the content of the journal is not constrained by a remit to appeal to or reach the interested general public. Whilst the papers must speak to the needs of students who are relatively inexperienced in philosophy, they presuppose that their audience is actively engaged in philosophy. Second, the content is serious in its focus on the central areas of philosophy. The big or traditional questions of metaphysics, epistemology, and ethics will provide the journal's centre of gravity. The third way in which the philosophy is serious is through the scope, variety and depth of analysis that can be achieved by the accumulation of papers over time. Moreover, each paper is not simply an introduction to one of the main topics on A-level, IB or degree courses. Such papers will indeed have a role in the journal, but they will not be the only kind. Our contributors will be offering original papers based on their own research. The journal will be a forum for the kind of critical engagement and debate that characterise the practice of philosophy. The fourth way in which the philosophy is serious is in the contributors themselves. The vast bulk of the papers will be written by professional philosophers engaged in both research and teaching.

About the Editorial Board

Stephen Grant is a full-time lecturer in philosophy at Richmond upon Thames College. He has also taught at King's College London where he is completing his doctorate on the emotions. His main interests are in the emotions, ethics and political philosophy. He has published on the ontological argument.

Dr Paul Sheehy teaches philosophy at Richmond upon Thames College and King's College London. His main areas of interest are in metaphysics, political and moral philosophy and the philosophy of the social sciences. His doctoral thesis was on the ontological and moral status of social groups, and he has published papers on social groups, voting and explanation and realism.

Paul Sperring is head of the philosophy department at Richmond upon Thames College and an A-level examiner in philosophy. He completed his undergraduate and masters studies at The University of Warwick, studying both analytic and continental philosophy. He is currently working towards his PhD at Birbeck College. His research interests are metaphysics and the philosophy of mind.

David Papineau

[Confusions]

about Consciousness

Introduction

Consciousness has suddenly become an extremely fashionable topic in certain scientific circles. Many thinkers are now touting consciousness as the last unconquered region of science, and theorists from many different disciplines are racing to find a 'theory of consciousness' which will unlock this final secret of nature. I am suspicious about all this enthusiasm. I think that much of the brouhaha is generated by philosophical confusion. In the end, I fear, there is no special secret of consciousness, and no special key needed to unlock it.

Dualism and Materialism

Much of the confusion about consciousness is generated by lack of clarity on the issue of dualism. The majority of scientists who are caught up in the current excitement about consciousness studies would probably deny that they are dualists, if the question were put to them explicitly. But at the same time I think that many of them are closet dualists. They strive to resist the temptations of dualist thinking, but as soon as their guard drops they slip back into the old dualist ways. The very language in which they normally pose the problem of consciousness gives the game away. 'How can brain states "give rise" to conscious feelings?' 'How are conscious states "generated" by neural activity?' The way these questions are phrased makes it clear that consciousness is being viewed as something extra to the material brain, even if the official doctrine is to deny this.

To help make the point clear, let me move away from the mental realm for a moment, and consider two contrasting analogies from purely physical science, the theory of electromagnetism, and the theory of heat. These two theories work rather differently. Think of how they relate heat and electromagnetism to other physical processes. The theory of the electromagnetic field is a theory of an *extra* physical entity, of something additional to other physical goings-on, such as the movement of charged particles. The charged particles are one thing, and the field they produce something further. But the theory of heat does not explain heat in a similar way. Heat is not something extra to the kinetic energy of moving particles. Rather, talk of the heat in a body is just another way of referring to the kinetic energy of the particles in it. There aren't two entities here, the moving particles *and* the heat. It's not as if a 'heat field' arises when the particles move. Heat is nothing but the movement of the particles, described in other terms.

Now, which of these is the better model for the relation between conscious feelings and brain activity? That is, should we expect a successful 'theory of consciousness' to show us how certain brain activities *generate* certain extra entities, the conscious feelings, on the model of the electromagnetic field? Or should we rather expect such a theory to show us how conscious feelings *are nothing but* certain brain activities, described in other terms, on the model of heat. When Francis Crick, for example, says that consciousness is associated with 40-Hertz neuronal oscillations in the visual cortex, or indeed when any scientist equates consciousness with any feature of brain activity, are we to understand them as saying that some extra conscious field is generated by the brain activity, or rather that consciousness is nothing but that brain activity, described in other terms?

We can call a theory of the former kind a dualist theory, and a theory of the latter kind a materialist theory. I suspect that much work in 'consciousness studies' simply hasn't decided whether it is aiming at a dualist theory or a materialist theory. The indecision matters because it can lend such work an air of spurious excitement. This is because a dualist theory of consciousness, while it would certainly be exciting, is a highly

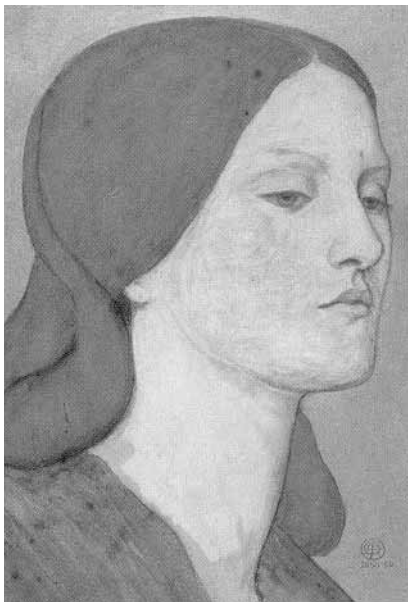

implausible prospect. A materialist theory, by contrast, while it is plausible enough, is not going to yield any exciting secrets. So, by fudging the issue between these two kinds of hypothesis, theorists of consciousness can have their cake and eat it. They can present their work as sharing the excitement of a dualist breakthrough, yet at the same time denying that its claims are any more surprising than a materialist hypothesis. If we are seriously to assess their theories, however, we need to be told whether they are intended in the dualist or materialist mode.

Consciousness and Life

To further clarify this issue, let me turn to another analogy, this time with the 'theory of life'. About a hundred and fifty years ago, scientists used to be excited about life in roughly the way that they are now excited about consciousness. While they were of course clear enough about which living systems are alive and which not, they were much perturbed by further questions. Why are those systems alive? What mysterious power animates them? And why is this power present in certain cases, such as trees and oysters, and not in others, like volcanoes and clouds?

These questions have now disappeared from active debate. Biology textbooks sometimes begin with a few perfunctory paragraphs about the distinguishing characteristics of their subject matter. But the nature of life is no longer a topic of serious theoretical controversy. Everybody now agrees that the difference between living and non-living systems is simply having a certain kind of physical organisation (roughly, we would now say, the kind of physical organisation which fosters survival and reproduction).

The reason for the nineteenth-century debate, and its subsequent disappearance, is that scientists used to be dualists about life, and aren't any longer. That is, they used to think that living systems are animated by the presence of a special substance, a vital spirit, or elan vital, which was postulated to account for those features of living systems, such as generation and development, which were thought to be beyond physical explanation. And of course, when they did believe in these vital forces, they then faced any number of exciting questions, such as why they arise in certain circumstances and not others, and what laws govern their operation.

But nobody is a dualist about life any longer. Nobody believes in vital spirits nowadays. A century and more of physiological research have persuaded scientists that the characteristic features of living systems can all in principle be accounted for in terms of normal physical forces, without bringing in any extra forces operating only in living bodies. With this realisation all the excitement about the nature of life has dissolved. To be alive is just to be a physical system of a certain general physical kind. There isn't any extra property present in living systems, over and above their physical features, which distinguishes them from non-living systems. So there are no pressing questions about the mysterious nature of this extra property.

I think that this story about life carries a direct moral for the study of consciousness. If you think that there are special mental forces, over and above the familiar physical forces, then you will think that there are exciting questions that must be answered, such as why these forces arise in certain circumstances and not others, and what laws govern their operation. On the other hand, if you think that the cognitive workings of intelligent beings depend on nothing but the operation of normal physical forces, without any extra forces operating only in brains, then you will see things differently. You may begin your textbooks with a few remarks about the distinguishing characteristics of conscious systems, but once this essentially classificatory question is out of the way, you won't want to spend any more time agonising about the nature of consciousness.

As my remarks so far will no doubt have intimated, I prefer the latter, materialist view of consciousness to the former, dualist story. And the reason is the same as in the case of life. All the physiological evidence indicates that no special mental forces are needed to account for the operation of intelligent organisms. Of course the evidence isn't conclusive, and doesn't absolutely prove that there are no such forces. But it weighs strongly against them. If you are unpersuaded, then ask yourself this question. Are any parts of matter in your brain ever caused to accelerate by mental causes, in the absence of any other forces? That is, do we need to include purely mental causes alongside gravity, the electroweak force, and so on, in the category of fundamental forces? As I said, there is no conclusive disproof of this thought, and it has its defenders, like Sir John Eccles (1989). But I take it that it would run counter to a large body of empirical evidence. (physicists would certainly be very interested if such a force could be shown to exist.)

Epiphenomenalism

No doubt some of you will be feeling uneasy about the analogy with life. Don't we have immediate access to the nature of conscious activity, via our introspective knowledge of our own minds, in a way that we don't have access to the nature of life? And doesn't this show us directly that conscious goings-on are distinct from any physical goings-on? When we are aware of a pain, say, or of seeing something red, don't we know directly that there is something going on in us which is quite different from any neuronal activity in our brains?

I agree that we all have strong intuitions to this effect. But they need to be handled with care. My own view is that they are illusory, and I will come to this in the next section. Still, some thinkers take these intuitions at face value, as showing that conscious feelings really are distinct from brain activity. However, if you take this line, then you face the argument from two paragraphs back, that physicists are going to be flabbergasted if it turns out that these extra conscious states sometimes cause bits of matter to accelerate in your brain.

There is one way of endorsing the intuition of distinctness without flouting physical orthodoxy and committing yourself to matter-accelerating conscious forces. You can insist that conscious goings-on are genuinely extra to brain processes, but deny them any causal powers. On this view, brain activity causes an extra conscious 'field', but this field then has no effect on brain activity. What happens in the brain itself is entirely accounted for by standard physical forces. The extra conscious field 'hovers' above the brain, but the brain runs along its own tracks, as directed by the standard laws of physics.

This position is known as epiphenomenalism. It too has its defenders. David Chalmers takes it seriously in his recent book, *The Conscious Mind* (1996). Chalmers attaches great weight to the intuition that conscious states must be distinct from physical states. Yet he is enough of a scientist to realise that it would fly counter to well-evidenced physical theory to credit these extra conscious states with powers to influence neuronal activity. So he suggests that perhaps they are just epiphenomenal 'danglers', caused by certain kinds of brain activity, but with no power to cause anything themselves.

Epiphenomenalism is a cogent position. But it has an obvious drawback. It forces us to deny that our conscious decisions are the causes of our bodily movements. It seems obvious that when I decide to raise my arm, or to go to the pub, my conscious decision is the cause of my limbs moving. But epiphenomenalists must deny this. On the epiphenomenalist view, our conscious mental life 'hovers above' the chains of physical causation that lead from my brain to my bodily movements, without playing any part in them. According to epiphenomenalism, we are like a child in a car with a toy steering wheel, happily twisting it this way and that, blissfully unaware that the actual movements of the car are quite independent of our attempts to steer it.

If we want to avoid this unhappy epihenomenalist picture, without positing Eccles-style extra mental forces, then we need to return to materialism. That is, we need to deny the intuition that the conscious states are separate from brain states, and insist that decisions and other conscious occurrences are nothing but brain activity, just as heat is nothing but molecular motion. Then, of course, we will have no difficulty understanding how decisions affect behaviour. For if conscious states are brain states, then all we need are the normal physical processes by which brain states cause behaviour. The puzzle about the causal role of conscious states dissolves. It is like the puzzle of how temperature changes manage to influence pressure, given that changes in mean kinetic energy already determine pressure changes. The answer, of course is that temperature *is* mean kinetic energy. Similarly, once we stop thinking in terms of two different states, conscious states and brain states, we don't need to tell any special story to find the conscious states some role in the causal processes. They already have one, in virtue of being one and the same as the brain states that cause behaviour.

Could conscious experiences really be one and the same as brain states? This seems perfectly coherent to me. This materialist position doesn't of course deny that it feels like something to be in some conscious state. It simply identifies this with what it feels like to be in some brain state. That is what it is like, for beings who are in that kind of brain state. (What would you expect it to be like to be in that brain state? To be like nothing? Why?)

The Antipathetic Fallacy

What about the direct intuition that brain states and feelings are quite different in kind? As I said above, I think this is an illusion. We are so close to our own feelings that it is easy to get confused about them. The trouble is that we are able to think about our feelings in a special way—by having them. And this special way of thinking about feelings makes it difficult for us to see that the things we are thinking about—namely, the feelings—are just the same things as we can think about in other ways—namely, brain states.

It will worth analysing this illusion of distinctness in a bit more detail, for I think that it is responsible for much confusion about consciousness. Let me assume at this stage that materialism is true. This doesn't beg the question. We have already in effect given a strong argument for materialism, by showing that the only alternatives are epiphenomenalism or Eccles-style extra mental forces.

The task still facing us is to explain why materialism should seem false, even if it is true. If we can explain this impression, on the assumption that materialism is true, then we will be home.

As a materialist, I maintain that conscious states are identical to brain states, just as heat is identical to molecular motion. But at the same time I recognise that we have two different ways of *thinking about* these states, two different kinds of concepts that refer to these states. By way of analogy, note how the everyday concept of *temperature* and scientific concept of *mean kinetic energy* both refer to the same quantity. Similarly, I say, with the everyday concept *pain,* say, and the physiological concept *nociceptive-specific neuronal activity.* These are two concepts that refer to the same thing.

There is something special about the mind-brain case, however, that makes it very difficult to accept that an everyday concept like *pain* can actually refer to the same thing as a brain state concept like *nociceptive-specific neuronal activity.* Note how mind-brain identity claims contrast with ordinary identity claims in this respect. Once we are shown the evidence, we have no special problem believing ordinary identity claims like temperature = mean kinetic energy. Not so with mind-brain identities. Even after we are shown the arguments for mind-brain identity, we still find it hard to accept that a conscious state can be one and the same as a brain state.

If you ask me, this is because concepts of conscious states pick out their references in a special way. In general, concepts refer to their objects by invoking some description. So when two concepts refer to one thing, this is normally because the two associated descriptions pick out the same thing. For example, The Evening Star and The Morning Star both pick out the same planet, Venus. Similarly temperature and mean kinetic energy can be regarded as two different description of the same quantity, one identifying it in terms of its macroscopic effects, the other in terms of its microscopic constitution. But mind-brain identities work differently. We have special ways of thinking about mental states—we refer to conscious states by using *imagination* or *attention* rather than description. That is, we imaginatively recreate the state in our mind, and then think of it as that *kind* of state (the kind I am recreating now). Alternatively, we focus attentively on the state while we are actually undergoing it, and again think of it as *that* kind of state (the kind I am attending to now).

When somebody refers to conscious states in these special ways, I shall say that they are exercising 'phenomenal concepts'. Such exercises require either that you are actually undergoing the experience referred to, or at least that you are recreating it in your imagination. Note how this means that uses of phenomenal concepts will share their 'what-it's-likeness' with the experiences they refer to. *People who deploy phenomenal concepts—that is, think directly about 'that experience'—will therewith have the feelings involved in the experience.* This is obvious in the case where they refer by attending to an experience while they are having it. But it also holds, to some extent, in the case where they refer to an experience by recreating it imaginatively. Visually imagining a red square is somewhat like actually seeing a red square. It isn't exactly like seeing, of course, but there is an obvious sense in which imagining and seeing are phenomenally similar from the subject's point of view. Similarly, an imagined pain shares some of the phenomenal unpleasantness of a real pain. It doesn't hurt as much, of course, or in the same way, but it can still make you feel queasy, or make you twitch, or make the hairs in your neck stand on end. In Hume's phrase, the imagining is 'a faint copy' of the original impression.

We are now in a position to explain why conscious states should seem intuitively so distinct from brain states, even if they are not. We are misled by the subjective commonality between uses of phenomenal concepts and the experiences thereby referred to. This subjective commonality can easily confuse us when we contemplate identities like *pains = nociceptive-specific neuronal activity.* We focus on the left-hand side, deploy our phenomenal concept of pain (that feeling), and therewith feel something akin to pain. Then we focus on the right-hand side, deploy our concept of *nociceptive-specific neurons,* and feel nothing (or at least nothing in the pain dimension—we may visually imagine axons and dendrites and so on). And so we conclude that the right hand side *leaves out* the feeling of pain itself, the unpleasant what-it's-likeness, and refers only to the distinct physical correlates of pain.

This line of thought is extremely common, both within philosophy and without. When we use our phenomenal concepts, we bring to mind, in a literal sense, something that feels like the experiential state we are thinking about. When we use non-phenomenal concepts, this does not occur. And this makes it seem to us that non- phenomenal concepts cannot possibly refer to the same experiential properties that are picked out by our phenomenal concepts. (Thus consider Colin McGinn's question, on the first page of his *The Problem of Consciousness* (1991), 'How can technicolour phenomenology arise from soggy grey matter?')

However, once we stop to examine it, we can see that this line of thought involves a simple fallacy, indeed a species of Quine's famous use-mention fallacy. There is indeed a sense in which non-phenomenal concepts (like *nociceptive-specific neuronal activity)* do 'leave out' the conscious experiences themselves. They do not *use* such experiences. But it does not follow that they do not *mention* such experiences. After all, most referring terms succeed in denoting their referents without using those referents in the process. There is no reason to suppose that non-phenomenal concepts of experience do not do this too.

Non- phenomenal concepts differ from phenomenal ones in not using the experiences they refer to. This is the sense in which they 'leave out' the experiences. But it does not follow that non-phenomenal concepts differ from phenomenal ones in what they mention. In this referential aspect, which is the one that matters, they need not 'leave out' any element of the experience, not even the 'what-it's-likeness'. There is no reason why we shouldn't be able to refer to this 'what-it's-likeness' using concepts which don't actually give us the feeling. It is only the peculiar fact that some special concepts, our phenomenal concepts, do refer by giving us the feelings which confuses us here.

In my *Thinking about Consciousness* (2002) I dubbed this confusion the 'antipathetic fallacy'. Ruskin coined the phrase 'pathetic fallacy' for the poetic figure of speech that attributes human feelings to nature ('the deep and gloomy wood', 'the shady sadness of a vale'). It seems to me that in the mind-brain case we commit a converse fallacy, and refuse to recognise that conscious feelings inhere in certain parts of nature, namely, the brains of conscious beings.

Conclusion

Let me sum up. If we want to avoid epiphenomenalism and Eccles-style dualism, we need to accept that conscious states are nothing but brain states, just as heat is nothing but molecular motion. Admittedly, this flies in the face of intuition ('How could technicolour phenomenology arise from soggy grey matter?') But intuition should not be trusted here. Even if materialism is true, we are easily seduced into thinking it false. This is because concepts of brain states ('soggy grey matter'), unlike phenomenal concepts ('technicolour phenomenology'), don't involve actual experiences. Still, this is no reason for thinking that the brain states themselves don't involve actual experiences.

Once we expose the antipathetic fallacy, then nothing remains in the way of accepting materialism. This will be good for the study of consciousness. But it will be bad for 'consciousness studies'. If we accept materialism, we will recognise that there are not going to be any breakthroughs, any crucial discoveries about what 'causes' consciousness. That would be like discovering what 'causes' life.

Of course there is no such thing to discover. All we can do is classify the different kinds of life, and try better to understand their mechanisms. Similarly with consciousness. We should stop getting excited about the spurious question of what 'causes' consciousness. Instead we should settle down to the serious business of classifying kinds of consciousness and exploring their mechanisms.

David Papineau

King's College London

References

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

John Eccles, *Evolution of the Brain: Creation of the Self. (*London: Routledge 1989.)

Colin McGinn, *The Problem of Consciousness.* (Oxford: Blackwell 1991.)

David Papineau, *Thinking about Consciousness.* (Oxford: Oxford University Press 2002)

Will Cartwright

John Stuart Mill on [Freedom] of discussion

In the West the freedom to say what you like, to criticise the authorities, and to discuss ideas openly and without fear, is agreed to be of fundamental importance. What there is less agreement about is when this freedom may be properly curtailed. It may seem surprising, then, that Mill devotes most of his famous account of these matters in *On Liberty*[1] to explaining why freedom of discussion is important and very little of it to what may seem the more pressing matter of when this freedom may be limited. But there are reasons for this apparent imbalance in Mill's discussion which are worth noticing.

Though freedom of discussion was widely accepted even in Mill's own day, he thinks that the arguments for it are not widely appreciated, something that is no doubt still true, and he holds, as he makes clear in chapter II, that one should not have beliefs without knowing the reasons for them. Moreover these arguments for free discussion have a wider relevance to issues of liberty, for Mill holds that these arguments, suitably adapted, are also arguments for freedom of action.[2]

There is a third reason for Mill's emphasis on the arguments for free discussion and for freedom generally. He thinks that freedom is increasingly threatened, not so much by the law as by an oppressive public opinion, in England at least.[3] Curbing this threat requires a widespread appreciation of why freedom of discussion and other freedoms are important. And the character of Mill's arguments for free discussion is instructive in this context. They do not particularly emphasise the predicament of those who have been forcibly silenced. Thus we do not hear about the peculiar frustrations of being prohibited from expressing one's view of the world and exploring it with others. Nor do we hear about the individual's right to free speech. As a utilitarian Mill rejects the idea of natural rights, and emphasises that society as a whole, not just the silenced individual, loses by the repression of free discussion. But this means that the social majority, which is the source of the oppressive public opinion that Mill fears, also loses by repression. And it is the case, as we will see, that Mill's detailed arguments emphasise that silencing people is in one way or another counterproductive, not just for society generally, but for the silencers in particular. One cannot help feeling that an important part of Mill's aim is to persuade those occasionally inclined to curb free discussion, who will surely sometimes include Mill's own civilised readers, and therefore us, that such curbs are self-defeating.

So Mill has his reasons for devoting so much of his discussion to the case for a freedom that we all ostensibly believe in already. The reason why, in contrast, he has so little to say about the limits to this freedom, dealing with them in the first paragraph of chapter III, is perhaps in part because his position on this is apparently very simple. However it is fair to add that Mill elaborates this issue, with reference to freedom more generally, later in *On Liberty*.[4]

Mill's case for liberty of discussion

Mill sets out his arguments for freedom of speech in chapter II of *On Liberty*. They are linked by a common concern with truth. The general idea is that truth is a casualty of the suppression of free discussion. In effect there are three arguments that are attached to three possible scenarios.

In the first you are to imagine that a majority who share a certain view seek to silence the minority who disagree. You are further to suppose that the majority view is false, as it happens, and the minority view is true. Mill argues that in these circumstances it is disastrous to silence the minority, disastrous for the *majority*, that is, because there is now no means of releasing it from its belief in a falsehood. If however the minority remains free to express its doubts about the majority view, then there is a chance that the majority will be brought to see the falsity of its view. This is a powerful argument.

The second scenario is the same as the first except that this time the majority view is true and the minority view false. Here a concern for truth might seem to support silencing the minority since its view is false. Suppressing falsehoods presumably supports truth. However Mill ingeniously denies this. He argues that if the majority silences its opponents, it will never have to defend its belief and over time will forget the arguments for it. But to have a belief without knowing the reasons for it is no way to hold a belief according to Mill. The belief may be true, but it is held as a prejudice. As well as losing its grasp of the arguments for its belief, Mill adds that the majority will in due course even lose a sense of the real meaning and substance of its belief. What earlier may have been a vital belief will be reduced in time to a series of phrases retained by rote. The belief will be held as a dead dogma rather than as a living truth. Finally, beliefs held like this are extremely vulnerable to serious opposition when it is eventually encountered. They are more likely to collapse because their supporters do not know how to defend them or even what they really mean. Mill thinks history repeatedly demonstrates this process at work and offers Christianity as an illustrative example. By Christianity Mill means the ethical core of the religion rather than its full apparatus of metaphysical beliefs, and he seems to think this ethical core is true. But by suppressing opposition to it over the centuries Christians have ironically weakened rather than strengthened Christian belief, and Mill thinks this explains the decline of Christianity in the modern world. Truth is, after all, a casualty of the suppression of falsehood.

Mill's third scenario involves both parties of opinion, majority and minority, having a portion of the truth but not the whole of it. He regards this as the most common of the three scenarios, and his argument here is very simple. To enlarge its grasp of the truth the majority must allow the minority to express its partially truthful view. These three scenarios exhaust for Mill the possible permutations on the distribution of truth, and he holds that in each case the search for truth is best served by allowing free discussion.

Assessment of Mill's case

The first and third of Mill's arguments are the most persuasive. If the majority view is wholly or partially false, then allowing critical discussion surely enhances the chances of truth replacing error. But if the majority view is already true, as in the second argument, allowing critical discussion does involve risk. In the rough and tumble of public debate people may be seduced away from truth to false ideas. What Mill does in this argument is to draw attention to the risks to truth involved in the opposite strategy of silencing criticism. But the balance of risks here makes this argument less persuasive than the other two.

This difficulty is compounded by the fact that majorities will always think that they are in the second scenario. They, like anyone else, will always take their beliefs to be true. Thinking a belief true is a condition of having it. So majorities will always take themselves to be in that one of the three scenarios where the case for allowing critical discussion is weakest. But Mill has another argument in chapter II which can be used to bolster his position at this point. He says that majorities, or anyone else for that matter, can never be entirely certain that their beliefs are true. They may always turn out to be false or partly false, in which case they will fall into scenarios 1 or 3 where the case for free discussion is powerful. Mill evidently supposes that this consideration ought to act as a powerful brake upon the temptation to silence dissent.

[Freedom]

Mill's defence of free discussion in terms of truth invites further questions on a number of points. First he has in mind discussion on a wide variety of issues, both theoretical and practical, such as scientific, religious, political and moral matters. But there is philosophical controversy about the applicability of the notion of truth to some of these issues. Many philosophers have held that moral judgements, and evaluative judgements generally, cannot be said to be true or false, but are in some sense or other inherently subjective. If this is correct, then free discussion about moral and other evaluative matters cannot be justified in terms of truth. A quite different justification will be required.

Secondly, is it really true, as Mill supposes, that we can never be certain that our beliefs are true? His position is that we can never have the degree of certainty that would warrant silencing criticism of our beliefs. That would be to claim infallibility for ourselves. If however we allow our beliefs to be criticised but no persuasive criticism is forthcoming, that gives us sufficient certainty to warrant our acting upon the beliefs, according to Mill. But this certainty precisely depends upon the belief continuing to be open to criticism. So the certainty requisite for action falls short of the certainty requisite for silencing criticism.

However consider my belief that 9x9=81 or my belief that I am writing this on a sunny day in the Suffolk countryside. Can I not be entirely certain that these beliefs are true? It is hard to believe that they could really turn out to be false, or that my certainty is only warranted if they remain open to criticism.[5] But even if in the case of these beliefs there is judged to be sufficient certainty to justify curtailment of discussion, what would be the point of such curtailment? These are not the kind of cases where people are tempted to silence criticism. When we come to the political, moral and religious cases where that temptation is real, then it is much more plausible to think, with Mill, that we cannot have certainty of the sort required to justify the curtailment of debate. For in these cases there are always considerations on both sides of the argument which have to be balanced. So it seems that when we have the sort of certainty that might warrant the suppression of criticism, we have no interest in such suppression; and when we have the interest, we lack the certainty.

But consider the following case where interest and, arguably, certainty come together. Suppose that the majority of Germans in the 1930s, observing the rise of Hitler and attending carefully to what he said and wrote, had concluded that his political vision was clearly false and very dangerous. This is arguably a case where the majority's certainty that it is right and he is wrong is of an order that would warrant silencing him. For it does not seem remotely credible to suppose that his view might turn out to be true and the majority's view turn out to be false. Silencing him might still be the wrong thing to do in these circumstances because it might make his views more attractive. But if it would be wrong, it would not be so because the majority could not be certain that it was right, as Mill supposes.

The limits of freedom of discussion

Mill uses an example to illustrate when free speech may properly be curbed.[6] He says that one ought to be free to attack corn dealers in the press as starvers of the poor, but that one should not be free to make the same attack orally to an excited mob outside a corn dealer's house. Even though the words used may be identical, the alteration of the circumstances in which they are uttered makes all the difference in Mill's view.

For in front of the excited crowd the words constitute 'a positive instigation to some mischievous act'.[7]

The principle embodied in this example is that one should be free to say what one likes unless one's saying it causes harm. This limitation on free speech, the only one acknowledged by Mill, commands very wide support, but only because the limitation as stated is so general. Once we get beyond this general formula into the more particular questions raised by it, agreement becomes much harder to reach.

First, what is harm? If the corn dealer is injured, or his house damaged, by the irate crowd, that is clearly harm. Rather differently, if someone's reputation is damaged by the propagation of falsehoods about him, that is also harm. It is the prevention of this harm which justifies the curtailment of free speech by the law of libel and slander. If we pass beyond these clear examples of harm to search for a more general account of it, it is plausible to understand it as a setback to the interests of a person. This means that understanding harm requires us to develop an account of what a person's interests are, which is no easy undertaking.

Secondly, how serious are the various sorts of harm? To curb free discussion to avoid harm involves sacrificing something of great value, and this can only be warranted if the harm to be avoided is grave enough. Mill makes it clear in chapter V, for this reason, that while harm is always a necessary condition of curbing liberty, it is not always a sufficient one. The harms threatened in the corn dealer case are at the graver end of the spectrum, but other harms are much less serious. All this means that we need, not merely an account of a person's interests, interference with which constitutes harm, but also an estimate of the value of those various interests.

Thirdly, how likely to occur must the harm be before it is proper to curtail free speech? Notice that in the case of the corn dealer the harms in question are very likely to occur and will be immediate. This is unlikely to be an accidental feature of Mill's example. It is indicative of the view that free speech is so valuable that it should not be limited to avoid merely speculative harm.

Mill's position on the limits of free speech is at once simple and complex. It is simple in so far as the avoidance of harm is the only limit he allows, but complex in so far as the application of this limit requires controversial judgements on a range of issues.

One thing that seems clear about Mill's position is that he does not think it proper to prevent words being uttered because people will be offended by them. One might regard offending people as a way of harming them or, even if one rejected this, one might still think that avoiding offence was an independent ground for curbing free speech, over and above the harm principle. But it is clear that Mill rejects both these options. Underlying this rejection, I think, is the judgement that since people are offended by so many things, curbing freedom of speech on this ground will deeply and unacceptably compromise that freedom.

Some examples will help to test out the claims of offence as a ground of interference. Consider the law of blasphemy which makes it a crime to attack Christianity in indecent or offensive terms likely to shock and outrage the feelings of Christians. The rationale for this intrusion into free speech is clearly the avoidance of offence, but there is widespread opposition to this rarely used law. Compare this with the crime of inciting racial hatred which involves the use of abusive or insulting words calculated to stir up racial hatred. Many of those who oppose the law of blasphemy favour this law. This is an inconsistent position if the purpose of this law is to stop people being offended by insulting attacks on them in virtue of their racial or ethnic character. If one is going to protect people from racial offence, why not from religious offence? However this overlooks the fact that the racial law speaks of inciting hatred, which the law of blasphemy does not, and one could argue that hatred is calculated to lead to harm in the form of racial attacks. That harm rather than offence is the rationale of the law receives support from the fact that the law on this matter is part of the Public Order Act 1986. This way of justifying the law would make the common position, of support for it and opposition to the blasphemy law, consistent.

[Freedom]

The two examples just discussed might be read as suggesting that Mill was right to reject offence as a ground of interference. But consider now other examples which might be thought to suggest otherwise.

In 1988 the writer Salman Rushdie published a novel called *The Satanic Verses* in which some of the characters think, say and dream things about Islam which caused immense offence to Muslims. The repercussions were extraordinary and global. There were demonstrations in Britain and across the world; the book was burned in public; demands were made for the banning of the novel; bookshops selling it were firebombed; some of those associated with the publication of the novel, and its translation into other languages, were assassinated; and finally, and famously, the Iranian religious authorities placed Rushdie under a fatwa which required devout Muslims to try to kill him. This unprecedented literary furore divided non-Muslim opinion in this country. Some thought it of overriding importance to reassert the freedom of the artist to express himself. Others were impressed by the degree of hurt in the Muslim community and thought that even artists must accept some constraints on their freedom of expression.

In exploring the latter view the two earlier examples become relevant again. During the controversy some Muslims applied to the courts to have Rushdie charged with the crime of blasphemy. The problem for the applicants was that the law of blasphemy only applies to Christianity and the courts decided that they did not have the power to extend it to Islam. Only Parliament could do that. As explained above, this way of limiting, or trying to limit, Rushdie's freedom to express himself in effect endorses offence as a limit on that freedom. But there is an alternative way of limiting his freedom which does not have to rest on offence. We could make it a crime to incite people to religious hatred, to parallel the existing crime of inciting to racial hatred. And the justification for this new crime, as for the existing one, would be that hatred is liable to lead to people getting harmed. Thus those who think that the Rushdie case shows the need for some restriction on freedom of expression can argue for it without abandoning Mill's position. They can argue for it in terms of harm rather than offence. But many will think this a false basis for the restriction and that the real issue is the offence to Muslims. They will hold that when the offence is this wide and deep and has to do with the ridiculing of a person's basic beliefs, then freedom of expression can properly be limited. This is, of course, to declare Mill's position inadequate.

A further, and interestingly different, example is that it is illegal in Germany and elsewhere to deny the occurrence of the Holocaust, the systematic attempt by the Nazis to eliminate the Jews and others.[8] Whereas the three previous cases involve insulting attacks on a person's basic beliefs or racial character, this one has to do with simply denying the occurrence of an event. Despite the difference the natural justification for this restriction is that denial of the Holocaust, by Germans in particular, is deeply offensive to the survivors of the event and their descendants. There are perhaps other considerations involved in this case. The identity of modern Germany is inextricably connected with the moral catastrophe from which it emerged. It defines itself against that calamity. For Germans to deny the Holocaust is to deny the central event in that calamity and, in a way, is to subvert the moral identity of the new Germany. Whether these further considerations are adequate reason for interfering with free expression is another question.

As a final example, consider freedom of expression with respect to pornography. In 1977 the government established a committee under the chairmanship of the philosopher Bernard Williams, to investigate what constraints there should be on the depiction of sex and violence in various media. Its report is instructive.[9] The committee thought that its remit raised fundamental issues concerning freedom of expression, and that in consequence it had initially to decide the grounds on which this freedom could be properly limited. First the committee accepted the harm principle. If pornography causes harm, then society can properly ban it. However, after reviewing the evidence on this controversial question, the committee concluded that it has not been proven that pornography does cause harm, though equally they thought that it has not been disproven. Secondly, despite the well-known liberal inclinations of its chairman, the committee went further and endorsed offence as a ground of interference. If there is controversy about whether pornography causes harm, there is none about whether it offends people – it clearly offends many. But, the committee thought, this does not entitle society to ban pornography, but only to constrain it in order to avoid offence. So the committee recommended that pornographic magazines should be sold in specially designated shops rather than being available in newsagents where they will give offence; that pornographic films be given explicit ratings that will warn

off the sensitive, and so on. By such means the committee sought to spare people offence while at the same time maintaining the freedom to produce and consume pornography, thus minimising the intrusion into freedom of expression.

After reflection on this array of cases, some will conclude that Mill's position is too minimal and that offence should be admitted as a constraint on free expression, in some cases at least. They can point out, as the Williams committee did, that the avoidance of offence need not involve the curtailing of free expression since it can often be accomplished by advance warnings to the sensitive about the nature of the material. Others will draw the contrary conclusion, with Mill, that the admission of offence will be too damaging to free expression. They may think it preferable to approach the problem from the other end, so that, rather than curbing free expression, we try to create a climate in which people are less susceptible to offence. At one time Christians in this country would have been outraged by the mere denial that God exists. Indeed there was a time when such denial counted as blasphemy. Happily both Christians and the law are now different. We should encourage similar developments elsewhere, it may be said. This is an important point, but it has its limits. A society entirely beyond shock and offence is one that we have good reason not to want, because such a society probably cares about very little.

Will Cartwright

University of Essex

Notes

1 Published in 1859; see especially ch II. The views of Mill discussed in this essay are to be found in ch II unless otherwise indicated.

2 Ch I last paragraph and ch III first paragraph.

3 Ch I.

4 Ch IV and V.

5 In the case of the second belief the problem of epistemological scepticism needs to be addressed.

6 Ch III first paragraph.

7 Ibid.

8 Amongst the other countries where it is illegal are France and Austria. A proposal from Brussels to make it illegal in all countries of the European Union is being opposed by Britain, despite the fact that in 1997 Mr Blair said there was a 'very strong case' for a Holocaust-denial law in Britain.

9 *Obscenity and Film Censorship. An Abridgement of the Williams Report.* (Cambridge University Press 1981.)

[Freedom]

John Collins

[Language]

A Dialogue[1]

Phil: Hi Lyn. I was hoping I would bump into you. I've heard that you have some quite bizarre ideas about language.

Lyn: Erm, that's rich. Are you still worrying about how you know you're not a brain in a vat?

P: Touché. Seriously, though, I've heard that you think that languages are mental states, properties of human brains - or 'mind/brains', as you say. Well, the thing is, Lyn, while we philosophers are notorious for our disagreements, if we agree about anything, it is that languages are not internal states.

L: Yes, it is indeed my view that 'languages' are states of the human mind/brain. I know that this is not the received view in philosophy, but it is none the worse for that. First off, instead of bothering about what the word 'language' means, let us consider some facts. The human brain is equipped with some resources - whatever they might be - which enable us to acquire any language - English, French, Latin, etc. This strikes me as obvious, just a fact. Well, if this is so, then the human brain has some property P that distinguishes us from chimpanzees, rabbits and rocks. P is innate, unlearnt - it comes with being human, just as flying comes with being a bird, or swimming comes with being a shark. Secondly,...

P: Hang on, hang on. I don't know what this 'P' is supposed to be, still less why it has to be linguistic in some sense. P might just be 'intelligence' or 'pattern recognition' or something at least which doesn't presuppose anything specifically linguistic.

L: You are perfectly right: P is just a stand-in for whatever it is about the human brain that distinguishes it from the brains of chimpanzees, rabbits, and everything else in the known universe. I didn't assume that it was 'linguistic'. Let us see if we can agree on something. Take a Japanese baby just born in Tokyo. Move on five years, what language will it be speaking?

P: Well, it depends, but I suppose you want me to say 'Japanese'.

L: You're right; we can't simply assume that Japanese babies will end up speaking Japanese; it depends on what language their parents and peers speak. That's my point! The typical Japanese child ends up speaking Japanese, but this particularity is not determined by the make-up of its brain. I mean, move the Japanese child to any other country and it will pick up the language of its new environment as easily as it would have picked up Japanese. A chimpanzee won't, nor will a rock.

P: I see what you are getting at. A child can acquire any language of any environment with equal ease, so whatever the child has innately - your P - can't correspond to any particular language; if it did, then acquiring one language might be harder than acquiring another, or even impossible, but that looks not to be the case. On the other hand, P just can't be any old property, for only humans have it; it must be able to cover the apparently indefinite variation in the world's languages, because any child can acquire any language. P is this capacity to target any language. Is this what you are saying?

L: Yes, exactly.

P: But now I'm confused. I thought that you want to argue that languages are internal mental states. But such a thesis doesn't follow from the thesis that there is some P. P corresponds to no language; we just showed that. P is merely the capacity to acquire a language.

L: You're jumping ahead. All we have established so far is that there is something about the human brain that enables the child to acquire any given language. This property we call P. Now comes a separate question: What is P? Is it something specifically linguistic, or is it some general capacity, which might be involved in doing mathematics or playing chess, or whatever? Well, my thought is that P is specifically linguistic without corresponding to any particular language; it expresses, we might say, the general or universal form of

languages, what they all share in virtue of being human languages.

P: Okay, let me see if I've got the idea. To acquire a given language we depend on some innate capacity, but, whatever language we do acquire, the same capacity would have worked equally well for any other language. Now, you also want to say that this capacity is specific to language, it's not some general purpose intelligence we employ outside of the realm of language. In a sense, then, this capacity is kind of like the essence of language - that which is invariant over all languages, that which underlies our capacity to acquire any given language.

L: Spot on. I call this fecund state of mind, Universal Grammar, or UG for short. You can think of it as the initial state of the *language faculty*, where later states of the faculty correspond to a speaker's acquired language.

P: But I still don't see what argument you have for this. I mean, why should there be any essence to language? Perhaps this P - the capacity to acquire any language - just does correspond to 'intelligence' or some sophisticated pattern recognition capacity. In other words, your UG is not a grammar at all; it's just the general capacity which makes us humans smarter than everything else in the known universe.[2]

L: Well, there might not be UG in the sense I mean; I have no mathematical proof that there must be such a thing, but I do have a very good argument. The child's linguistic experience - the sentences she hears, etc. - is not rich enough for her to figure out how the language as a whole works. In other words, there's just a real gulf between what the typical child experiences and what she ends up knowing: this complex thing we call her language. She requires additional information, over and above the mere capacity to organise and segment her experiences, in order to start creatively using language herself, rather than merely repeating what she has heard. Well, if the required information doesn't come from outside, it must come from 'inside', namely, UG. UG expresses what the child brings to language as part of its biological make-up.

P: Erm, this is style of argument goes back to Plato's dialogue *Meno*, where Socrates shows that Meno's slave boy understands Pythagoras' theorem, even though no-one has taught him any maths. Plato thought that this meant that the slave boy's soul brought the mathematical knowledge with it from whence it came - the realm of the Forms. You've just replaced Forms and souls for biology. More generally, you are siding with the Rationalists - Descartes and Leibniz - against the empiricists, like Locke and Hume. You think that we have innate linguistic ideas, concepts which don't derive from our senses. We are not 'blank slates'.

L: That's right, I'm a Rationalist in modern dress. It's worth pointing out, though, that no-one has ever really believed that the mind was entirely blank at birth. I mean, it just makes no sense. If it were blank, how could it be that every human ends up speaking a language but our pets don't? Everyone thinks that something is innate - unlearnt - I differ in thinking that something specifically linguistic is innate.

P: Okay, I take your point: if my mind was entirely blank at birth, then I suppose, initially, there wouldn't be any difference between me and my dog, but we end up differing greatly. Your point is: We only end up so different because we were very different to begin with. I can see this. But I still don't see why what I start with, as it were, has to be linguistic in some sense. Why can't I just be a lot smarter than the dog?

L: As I've said, because the experiences you have just aren't rich enough for you to arrive at, say, English on the basis of just being really smart, whatever that might mean.

P: Surely the child is inundated with language. Anyhow, this is an empirical issue for psychology. What basis do you have for assuming that the child lacks this or that information about language from what it hears around it?

L: You misunderstand my point. I haven't made any assumptions about the linguistic data the child is likely to meet. You are quite right. I really don't know how rich the data are. But my argument is that no matter how rich they are, they aren't going to be rich enough. In linguistics, we have as much data as we could hope for, but we still can't figure out the true generalisations for English. Linguists are very bright, highly educated people, working together with a rich tradition of research behind them. Even so, they can't figure out how English, or any other language, works. We're trying, for sure, but without clear success yet. So, give the child - minus UG - as much data as you want. No, give the child the acquired wisdom of 2000 years of thought about language, and it still won't be able to figure out what language it is supposed to be speaking. There is a paradox here. The linguist relies on data - lots and lots of data, as much data as she needs - but she can't figure out how English works. On the other hand, the child picks things up by about the age of five, making remarkably few mistakes along the way, and independent of other variables, such as intelligence, social background, disability, etc.[3] This is miraculous, yes? Well, no. It's only miraculous if we make the assumption that the child picks the language up from the environment without something like UG. The child's task is easy precisely because it doesn't rely on data; most of what it needs to know is already built into its brain; the child just waits for it to be triggered. The environment plays a part here, but only to *select* options already available. It doesn't *instruct* the process. This is wildly picturesque, but it's as if the child already knows English, French, etc., and after hearing certain kinds of very simple patterns that uniquely cluster with French, the child's brain selects French, rather than Swahili, say, which exhibits different kinds of simple patterns. So, the linguist's task is really hard precisely because she must rely on data, whereas, in effect, what she is trying to understand is what the child has innately without relying on data, namely UG.

P: I'm not sure I understand what you're saying. The idea that we have every possible language in our heads at birth is just silly.

L: I don't know, perhaps we do: these matters are empirical. Anyhow, we don't have to take the picture seriously.[4] A common idea is that a speaker has a menu of fixed options and some broad constraints about what can go with what. It's as if there's a restaurant in your head, and you can have the veal or the trout, but if you have the veal, you can't have the salad, you must have the soup. This gives you two meals. Add to the menu, and you have a whole lot of other possible meals. Each language, or future state of the language faculty, corresponds to some selection of options, a meal, as it were. Experience selects the options - orders the meal - but it doesn't create the options, nor dictates what options can go together. The chef is in your head; experience is the customer. So, French need not be sitting there in the child's brain twiddling its thumbs waiting to be called; the language faculty doesn't realise all of its possible states simultaneously, just as not every meal is waiting in the kitchen of a restaurant - we're not talking about McDonalds here.

P: Okay, I think I see what you mean, but this idea is very strange. You will admit that it is highly counter-intuitive.

L: Most science and philosophy is counter-intuitive. Perhaps the problem for intuition in the present case is that it is an example of a general phenomenon: we know an awful lot we don't know we know! Most of our knowledge about language is *implicit*, unconscious. What we have found is that there is a whole range of significant generalisations about the structure and meaning of English sentences - let's just stick with English, but the point applies generally - that we clearly understand and follow, but which we don't know consciously. Further, these generalisations can't be stated in terms of the visible properties of sentences, as it were. These are properties of order and sound. In short, we know things that we couldn't have acquired just from the data. Linguists have only just noticed many of them. In some sense, they come from UG.

P: Could you give me some examples?

L: Sure. Everywhere you look, you find the same thing. Consider verbs to do with liquid, like 'spill', 'squirt', 'pour', etc. Just to take 'spill', although the

example generalises, we can say both (2) and (3):

(2) Bill spilt water on the floor

(3) Water spilt on the floor.

(2) implies (3). We can all recognise this. But now think about verbs to do with substances of a certain viscosity, such as 'smear', 'daub', 'plaster', etc. Again, just to pick one, we find a different pattern from that above:

(4) Bill smeared paint on the wall.

(5) Paint smeared on the wall.

It is not so much that (4) doesn't imply (5); rather, (5) is just ill-formed; it is not an English sentence.[5] Well, why can't paint smear on the wall in the way that water can spill on the floor? It's no good thinking about the difference between paint and water. I mean, no chemist is going to do research on this question. That's just daft. More to the point, no-one teaches this to children acquiring English. Indeed, how could it be taught to children when adults don't even know it explicitly? Nevertheless, it's part of being a competent English speaker that one finds no problem with (2)-(4) but finds (5) distinctly odd. Well, no-one ever utters (5); it's just not part of English.

P: Erm, I had never noticed the difference. But I don't see the big mystery. Let me think... the difference is just that verbs like 'spill', 'squirt', etc. don't require someone to do the squirting, spilling, etc., whereas verbs like 'smear', 'daub', etc. require someone to do the smearing, daubing, etc. That seems simple enough. We work out the difference via knowing what the verbs mean.

L: Remember, we are thinking about the child acquiring the language. It is perfectly correct to say that 'spill', in some sense, doesn't imply an agent who does the spilling, whereas 'smear' does imply an agent who does the smearing. But where does this difference come from? It seems to me that this difference is the very same difference I initially pointed out - you haven't explained anything. Besides, how does the child learn what the words mean? Do you imagine that we sit our children down and tell them that 'spill' doesn't involve a notion of agency whereas 'smear' does? This obviously doesn't happen. Nor is it remotely plausible to think that the child can pick the difference up from its observation of liquids. I've never seen water spill itself on the floor - someone does the spilling - just as I've never seen butter smear itself on my toast in the morning.

P: Okay, okay. I'll have to think about this a lot more. I'm sure you have lots of other examples...

L: Consider double-object constructions... [6]

P: Hold on! What I was going to say is that I still don't see why you think that languages are internal mental states. It seems to me that all this talk about English presupposes a shared public language. Have I missed something?

L: Well, English is a 'convenient fiction'.

P: What?

L: Look, you've granted me UG; or at least you have granted me that there is some property of the human brain that distinguishes us from chimpanzees and the rest of the universe as regards language. Well, is this property just some general capacity or is it specific to language? It seems to be quite specific. We are not taught the subtle difference between 'spill' and 'smear', but we all acquire the words fine. This knowledge must come from UG, in some sense, for it doesn't come from outside. But UG is just a property of the brain. In other words, the states of our language faculties are just variations on the menu provided by UG.

If the brain develops in this way, you get 'English'; if in that way, 'French'; if in some other way, 'Latin'. I use the 'scare quotes' because 'English', 'French', etc. are just picking out a set of language faculties that are sufficiently similar to support various generalisations (e.g., English sentences require subjects, Italian ones don't), although each will be quite distinct in all sorts of ways. For instance, we tend to speak of 'Chinese', but there is much more variation between Chinese dialects then there is between French, Italian, Spanish and Romanian, say. Languages are a bit like races: Europeans, Africans, etc. Races aren't real as far as biology is concerned; that is, no interesting generalisation in biology concerns Africans or Europeans. Biology is concerned with interbreeding populations which cross-classify what we think of as races.

P: So, you're saying that really we each have our peculiar way of speaking and understanding, our individual dialects if you will, which can be more or less similar within the bounds set by UG. Talk of English is a crude classification, much like our talk of Africans.

L: Yes, that's right. There's a joke in linguistics that a language is something with an army and a navy. For all kinds of reasons to do with politics, culture, history, whatever, we are interested in racial differences, much to the detriment of human development, I might add. Just so, for similar reasons we are interested in classifying people as English speakers or French speakers, but this is interest-relative, as they say: there is no thing - English - which all and only those we want to call English speakers know.

P: Oh, look. Just because the notion of English is a bit vague around the edges, it doesn't mean that it doesn't exist. It's a social object, like our constitution, or the welfare state. Our dialects are not internal states!

L: You misunderstand me. I'm not saying there is no such thing as English. I'm not even sure what that means. All I'm saying is that English is not a notion which enters into serious investigation of language, apart from as a useful term to talk about a particular population of speakers: people of the British mainland (and scattered former dominions), or something like that. The linguist is interested in the mental states of the speakers, not some object - English - which they all relate to in some mysterious way. We have granted, I hope, that the speaker must have some very complicated mental structures to be able to speak and understand in the normal way. The character and development of these structures is what interests me. Anything else you want to say about language will presuppose such structures. Of course, dialects are not internal states, but they depend on specific internal states, states which distinguish us from one another, but still mark us out as language users in distinction to chimpanzees or rocks. What's the problem?

P: Well, look, we have a shared store of thoughts. If I manage to communicate with you, then we have both grasped the same thought; I don't have to infer what you mean - I can just understand your words, because they express the same thought which I would express were I to utter those words. Any worthwhile conception of language must make sense of its prime function - communication. Your conception makes communication a mystery.

L: This talk of grasping thoughts is just a metaphor, which, to be frank, I don't understand. Besides, I really don't think that language is *for* communication. I can communicate with language, but birds and bees can also communicate without language. If you want to say that language just is communication, then there is nothing much else to say. The concept of 'communication' is a bit like 'swimming' in this regard. Lots of things swim: humans, sharks, otters, hippos,... they propel themselves through water. There's really nothing much else to say about swimming in general; although a particular account of how a shark swims will be of interest to a marine biologist, it won't tell you anything in particular about how humans swim.

P: I'm sorry, you're being a bit oblique.

L: Look, classifying creatures as swimmers is just not interesting, it's merely descriptive. Just so for classifying creatures as communicators. There are a multitude of ways to swim and communicate which our familiar notions just don't discriminate. So, sure, language allows for communication, but so does having legs, or a big tail, or scent pouches, or a changeable skin pigment. None of this has much to do with human language in particular.[7]

P: Okay, communication might not be an interesting notion in itself, but your position appears to make it impossible.

L: Let's back-track. We both agree that the child, and so the adult, requires quite sophisticated mental structures to use language, including, presumably, communication. Y'know, our heads aren't just full of baked beans; there is a lot of dedicated, very complex machinery in there. No-one tells us what 'spill' means, but it turns out to be quite complex; that comes from inside. Okay, well, the question is: Do we need more than these mental structures to understand

communication? As far as I can see, you simply assume that we do, but I just don't see it. Imagine Smith and Jones having a conversation. On my view, this is possible because of the way things are with Smith and Jones, period. Neither has to grasp any independent thing. If their internal states are sufficiently alike (e.g., 'spill' doesn't involve agency, but 'smear' does), then they will respond to each other in a coherent way. This is not an all or nothing thing. Often, understanding just breaks down - a conversation is a precarious affair. On your view, on the other hand, to make sense of Smith and Jones we require not only a description of the respective mental states, but also a specification of some third thing: a shared language which expresses these strange things we are supposed to grasp. The burden is on you to say why.

P: Erm, it's getting late and I have to write a paper about brains in vats. I'll have to get back to you on this. It is all quite intriguing, but it's, er,... how should I put it?... unattractive. A bit like someone not so long ago, you seem to be saying, 'There's no such thing as society. There's just you and your brain'.

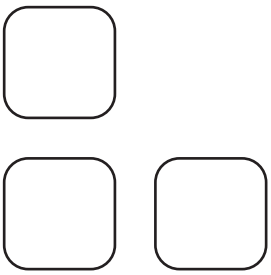

L: Politics has really nothing to do with it, at least not explicitly: no political position - either reactionary or progressive - is deducible from what I've argued. That said, political and moral views do presuppose, albeit often inchoately, some conception of human nature, much as, say, any theory of education presupposes some conception of the child's natural capacities and motivations - we are not dealing with chimpanzees or rocks. Well, my view of our nature is that we are not blank, malleable beings to be shaped and molded according to whatever structures of education and socialisation are current. We each bring with us an individual creativity with language, on a biological theme, from which we spin our very distinct perspectives on the world and each other. If anything follows from this, it is that we should treat each other with respect and dignity, and seek to have such attitudes reflected in social organisation, for we are not mere products of society. Otherwise put, a rich conception of human nature, in contrast to the poor one of the standard empiricist position, places the burden of justification on those who would seek to wield power over us - socially, educationally, politically, personally, whatever - in the name of some higher goal or authority, such as the state, the nation, the party, the race... Mind, we shouldn't need philosophy or science to tell us this.[8]

P: Well, thanks for the chat, Lyn - every time I spill my drink, I'll think of you.[9]

John Collins

University of East Anglia

Notes

1 Lyn expresses the views of Noam Chomsky. Chomsky's work on language is voluminous and sometimes highly technical, but always clear. A good starting point is *Rules and Representations* (New York: Columbia University Press, 1980) or the first two chapters of *Knowledge of Language: Its Nature, Origin, and Use* (Westport: Praeger, 1986). He is often at his best in interview; see *Language and Politics*, ed., Carlos Otero (London: Black Rose Books, 1988). For a general introduction to Chomsky's work, see Neil Smith, *Chomsky: Ideas and Ideals* (Cambridge: Cambridge University Press, 1999). For a wonderfully accessible introduction to the current 'Chomskyan' framework in linguistics, see David Adger, *Core Syntax: A Minimalist Approach* (Oxford: Oxford University Press, 2003).

2 Phil's complaint here has been often expressed; for the classic example, see Hilary Putnam, 'The 'Innateness Hypothesis' and Explanatory Models in Linguistics', in his *Mind, Language and Reality* (Cambridge: Cambridge University Press, 1975), 107-116. For a recent example, see Fiona Cowie, *What's Within: Nativism Reconsidered* (Oxford: Oxford University Press, 1999).

3 For an overview of the empirical research in this area, see Steven Pinker, *The Language Instinct* (London: Penguin Books, 1997).

4 For a model on which, in a certain sense, all possible language are 'in our heads', see Charles Yang, *Knowledge and Learning in Natural Language* (Oxford: Oxford University Press, 2002). The book is not for beginners, although the opening chapters are accessible.

5 (5) can be *part* of a sentence, as in 'Bill saw paint smeared on the wall', but 'paint smeared on the wall' is here a noun phrase, with 'smeared' as participle, it is not a sentence in its own right.

6 Lyn was going to point out that whereas 'Bill told Mary the message' is well formed, 'Bill reported Mary the message' is not. This is so, even though 'tell' and 'report' are near synonyms, and both admit the non-double-object construction: 'Bill told/reported the message to Mary'. Again, no-one tells us this.

7 For a survey of systems of animal communication, without serious mention of human language, see Marc Hauser, *The Evolution of Communication* (Cambridge, MA: MIT Press, 1996). For a speculative proposal on the phylogenic relation between language and animal systems of communication, see Marc Huaser, Noam Chomsky and W. Tecumseh Fitch, 'The Faculty of Language: What is it, Who has it, and How did it Evolve?' *Science*, 298, November 2002, 1569-1579.

8 Chomsky is perhaps better known as a political/human rights campaigner than as a linguist/philosopher. He is, however, quite reticent to speak about the connections between these two strands of his thought. For the connections that do exist, see Noam Chomsky, *For Reasons of State* (London: Fontana, 1973) and *Chomsky on Education and Democracy*, ed. Carlos Otero (London: Routledge, 2003).

9 My thanks go to Noam Chomsky for inspiration and helpful comments.

Paul Sperring

Thinking about [Zombies]

I Zombies

A number of contemporary philosophers have made use of the idea of zombies to defend a certain position within the philosophy of mind. What they have in mind are creatures very distinct from the sorts of zombies imagined by horror-filmmakers or Haitian occultists. *Philosophical* zombies are just like you and me in many respects. In fact, they are exactly like you and me in all but one important respect. They are physically, behaviourally and functionally identical to us but utterly without any conscious states.

Let us imagine that I have a zombie counterpart. We will call him $Paul_z$. Let us further imagine that this counterpart inhabits a world pretty much like the actual world. We'll call it $world_z$ (or W_z). This counterpart will, in W_z, do and say whatever I do and say here in the world. Let us imagine, then, that $Paul_z$ is currently tapping away at a word-processor writing a paper about zombies in W_z.

If we were witness to $Paul_z$'s actions then we would be able to discern no difference between what he is doing there and what I am doing here. All goings on in W_z, from the outside, would appear to be identical to the goings on in the actual world.

My experiences, tapping away at the keys, will have features, however, that $Paul_z$'s will lack. Properly speaking, $Paul_z$ will lack *phenomenal* experiences entirely. I am currently having a variety of phenomenal experiences, each with a distinctive qualitative feel: the texture of the keys under my fingers; the tapping and clicking noises made by my fingers striking the keys; the varied colour experiences of the computer monitor and objects surrounding it on my desk; the smell and taste of the tea just now sipped. All of these things will be denied to $Paul_z$ in his phenomenally textureless, noiseless, colourless, odourless and tasteless W_z. Of course $Paul_z$ will appear to have all of these experiences, and when someone asks about his cup of tea he will say 'very nice, thanks, just right, not too hot', or something like it.

It will appear, from the outside, that $Paul_z$ is having all sorts of conscious experiences. An observer, on being told that $Paul_z$ had no conscious experiences at all, might, of course, be puzzled about how it was that $Paul_z$ could type meaningful sentences employing references to phenomenal experiences unless he had had such experiences. However, we will put this worry aside here. Let us merely reiterate that the zombie counterpart has no phenomenal experiences at all, despite appearances.

II Attacking the Identity Thesis

So, this is what the philosopher has in mind when she introduces the idea of a zombie into her discourse. What purpose is served by thinking about zombies? They are introduced, usually, as a means of showing that the identity thesis about mind and body is false. The argument runs as follows:

- If *a* is identical to *b* then *a* is necessarily identical to *b*.
- If *a* and *b* are identical then there is no possible world where there is *a* but not *b*.
- There is a possible world where there is *a* but not *b*.
- Therefore, *a* is not identical to *b*.

To cash this out we can see that in W_z there are functioning brain states of a certain sort[1] (cases of *a*) but no conscious states (cases of *b*), as with $Paul_z$. But if conscious states were identical to brain states, hence necessarily identical, then there couldn't be a world where there were functioning brain states of a certain sort without conscious states. But we know that it is conceivable that there is a world where brain states occur without conscious states – we have just conceived of the world above (W_z), where $Paul_z$ is typing – and since whatever is conceivable is possible it must follow that brain states and conscious states are not identical.

To make the argument transparent let us run it again, attempting to identify particular brain states (Z-fibre firings[2]) with particular conscious states (phenomenal experiences):

- If Z-fibre firings are identical to phenomenal experiences then Z-fibre firings are necessarily identical to phenomenal experiences.
- If Z-fibre firings and phenomenal experiences are identical then there is no possible world where there are Z-fibre firings and no phenomenal experiences.
- In W_z there are Z-fibre firings but no phenomenal experiences.
- Therefore, Z-fibre firings are not identical to phenomenal experiences.

If this argument is sound then the identity thesis about mind and body is in trouble. Is it sound? I want to suggest a couple of lines of attack against the argument, but in its strongest versions it has been taken by a number of philosophers to be compelling.[3] One way of undermining the argument would be to deny that zombies are conceivable. If it could be shown that we cannot actually conceive of such things then we lack the grounds for claiming that they are possible - that is, W_z would not be a possible world. This would be an interesting strategy to adopt, but the defender of the zombie argument might be puzzled as to what it is that they have been thinking about all along. The objector will have to explain the apparent ease with which we think about zombies, that is, if the objection is that zombie thoughts are impossible thoughts then what is it that we have been conceiving of if not zombies? Ersatz zombies perhaps? I will say a little bit about this later, but first I want to explore another possible way of attacking the argument from zombie conceivability, *viz*., the denial that conceivability entails possibility.

III Does Conceivability Entail Possibility?

Despite its appeal many commentators have objected to the claim that whatever is conceivable is possible. The claim, and a forebear of something like the zombie argument, can be traced right back to Descartes who also helps himself to a refutation of materialism on the strength of it. Descartes' argument depends on the following principle.

(C➝P) If we can conceive of some state of affairs S then S is possible.

There is *prima facie* plausibility to this claim. It can easily be shown how conceivability is a pretty good guide to what is and isn't possible with some examples. We can conceive of Brighton and Hove Albion winning the FA cup this year, or we can conceive of a British rail company running all of its trains on time all of the time. Someone might object that these are unlikely scenarios, but it would be rather odd for that someone to claim that they are *not possible* states of affairs. On the other hand we cannot conceive of a square circle or a married bachelor, and we are evidently right to conclude that these things are not possible states of affairs. These latter are inconceivable for Descartes because they are self-contradictory, we simply cannot, for instance, think the thought of a person who, simultaneously, both has and does not have the property of being married, whereas there is nothing self-contradictory in thoughts of the former type involving football matches or efficient public transport systems. So conceivable states of affairs are states of affairs that are possible (or *not im*possible, which is the same thing). So if Descartes is right here then if we can conceive of a certain something then that thing is possible.[4]

It is on this basis that Descartes constructs his argument for dualism. He thinks that we can conceive of minds without bodies. The parallel with the zombie argument will be obvious, but instead of mindless bodies we get disembodied minds (ghosts, rather than zombies). The following is Descartes' argument as it appears in Meditation VI:

> First, I know that everything which I clearly and distinctly understand is capable of being created by God so as to correspond exactly with my understanding of it. Hence the fact that I can clearly and distinctly understand one thing apart from another is enough to make me certain that the two things are distinct, since they are capable of being separated, at least by God. The question of what kind of power is required to bring about such a separation does not affect the judgement that the two things are distinct. Thus, simply by knowing that I exist and seeing at the same time that absolutely nothing else belongs to my nature or essence except that I am a thinking thing, I can infer correctly that my essence consists solely in the fact that I am a thinking thing. It is true that I have [...] a body that is very closely joined to me.
>
> But nevertheless, on the one hand I have a clear and distinct idea of myself, in so far as I am simply a thinking, non-extended thing; and the other hand I have a distinct idea of body, in so far as this is simply an extended, non-thinking thing. And accordingly, it is certain that I am really distinct from my body, and can exist without it.[5]

To repeat, Descartes is accepting the claim here that whatever is conceivable is possible, but he is not appealing, as in the zombie argument, to possible worlds. Instead he employs God to make the conceivability entails possibility thought (hereafter C➛P) transparent. For any clear conception that we have of a state of affairs S God could make S obtain.[6] So if I have a clear and distinct idea of something then that thing just is possible. Once we accept this then we can see how the conclusion follows. If mind and body were identical then a clear and distinct conception of the essence of one would necessarily bring with it a conception of the other. But if we can conceive of states of affairs where there are minds but nothing physical or bodies but nothing mental then mind and body are possibly distinct. According to Descartes we can have a conception of mind that excludes any physical features and a conception of body that excludes any mental features so they cannot be identical. Hence dualism.

However, it might be objected, just because I can conceive of minds and bodies as separable does it really entail that they are possibly separable? Does conceivability in this instance entail possibility?

Antoine Arnauld constructed a counterexample that questioned Descartes' use of the C➛P principle.[7] In Arnauld's objection he asked us to imagine a confused geometer who understands some basic properties of triangles but who is not familiar with all of the properties of them. In particular the geometer does not know that all right-angled triangles have a certain property ϕ (the property expressed by Pythagoras' theorem). He then conceives of a right triangle T that lacks ϕ, and (by C➛P) concludes that it is possible that there is some such object (T without ϕ). However, T, if a right triangle, *necessarily* has property ϕ, so it is not possible that T without ϕ obtain. But if conceivability does not in all instances entail possibility then need we accept that the conceivability of disembodied minds entails that there could be such things?

Descartes, in answering Arnauld, points out that unlike the confused geometer he, Descartes, has a *clear and distinct* conception of mind and body. The geometer concludes that T without ϕ is possible precisely because he doesn't understand what essentially belongs to T and all right triangles. If indeed we supposed him to be an even worse geometer then there is no telling what sorts of things might, by his own lights, be included in the space of geometric possibilities – maybe triangles whose internal angles were greater than 180°, or perhaps even square circles. So it is not conceivability *simpliciter* that entails possibility but clear and distinct conceivability.[8] Thus we have a stronger version of $C \rightarrow P$.

> ($C_{cd} \rightarrow P$) If we can clearly and distinctly conceive of S then S is possible.

What does Descartes mean exactly by 'clear and distinct' conceivability? In Meditation III he spelt out its use as a criterion for certainty – anything that can be so conceived must be true. Thus anything that cannot be doubted, that must be assented to when thought about (the clear part), independent of elements that are doubtable (the distinct part), could not possibly be false. The *cogito* shows this clearly: I cannot doubt that I exist, it is impossible to think that I do not (I *clearly* conceive of myself, as a thinking thing, as existing), and I am thinking here only of my capacity to think, independent of (*distinct* from) thoughts about my body or some other thing.

So, having realised that he cannot but be certain about his own existence Descartes asserts that there must be something in the manner of his knowing this that can be generalised as a test for certainty. Anything else known in this manner will have the stamp of certainty.

There is, however, an immediate note of caution struck in Meditation III. Descartes accepts that some things might at first appear to be known with clarity and distinctness, but turn out, on more careful reflection, not to be. In a pre-philosophical state this is how things are with most of us – e.g. when we ordinarily assume that our senses give us reliable and accurate knowledge of the world as it really is. Having recognised that it is a problem that needs to be addressed, however, it is not at all clear that Descartes has a solution to it. And it is exactly this problem that gives Arnauld's objection its teeth.

Arnauld's attack is supposed to be defeated by $C_{cd} \rightarrow P$ because the confused geometer had no clear and distinct conceptual grasp of triangles and their necessary properties. He had an unclear, imprecise conception of what belonged to triangles and hence had no warrant for the claim that what he thought about them was genuinely possible. Now, Descartes' stronger version of $C \rightarrow P$ might work if we could pin down (a) what we mean by 'clear and distinct' conception and (b) when we have epistemic warrant for claiming that the conception under consideration counts as a veridical case. Unfortunately for Descartes he is unable to provide a satisfactory account of either. In the case of (a) – aside from appealing to some unsatisfactory characterisations, such as 'that which is manifest to the natural light of reason' – Descartes can only appeal to cases where we do grasp the thing in question clearly and distinctly as *illustrative* of the conditions required. The best (and perhaps only) illustrative example, as we saw above, is the *cogito* where we are just compelled to see its truth as soon as it is presented to us.[9] With respect to (b) this just brings us right back to Arnauld's objection. It seems right to say that the geometer *would* change his mind about thinking the triangle in question possible if his conception was much clearer. But what if the geometer in question had thought his conception to be clear and distinct. It does not seem enough to say that he was just wrong to think he conceived things clearly and distinctly – intuitively, it is obvious that he just does not see things aright – because Arnauld is looking for a warrant independent of the *seeming to be* in the grip of clarity and distinctness. In other words, for our discussion, what allows us to distinguish cases of *apparent* $C_{cd} \rightarrow P$ and cases of g*enuine* $C_{cd} \rightarrow P$? This goes right to the core of the question whether (and when) conceivability entails possibility. If we have no epistemic warrant then we are merely appealing to the gut, and if this was not enough for acceptance of $C \rightarrow P$ then it is difficult to see how it will suffice for $C_{cd} \rightarrow P$.

What one needs from a refined version of $C \rightarrow P$ then is some sort of internal guarantee that allows one to hold one's conceptions up to the light and see that they are genuine cases of conceivability that entail genuine possibilities. That is, we need something that will ensure that conceivability infallibly picks out possibilities.

Before considering whether there are any better versions of $C \rightarrow P$ that will do the work here I want to see if something like Arnauld's objection might be deployed against the zombie conceivability argument. Might the person who conceives of the zombie be like the confused geometer? It might be thought that the two cases are disanalogous since the geometer

was thinking of something that was logically impossible whereas there seem to be no logical or conceptual constraints on our thoughts about zombies. Better thinking *a priori* on the geometer's part would have revealed to him why there could not be a right triangle that lacked the ϕ property. Could better thinking *a priori* rule out zombies? What if we, committed to the truth of the identity thesis (perhaps having independent reasons for accepting it), just stipulated that if anything had property *a* it would of necessity have property *b* – and then it would be an analytic truth that there could be no *a*'s without *b*'s, i.e., no zombies, so they can be ruled out *a priori*. We could simply claim in advance that if pains simply were Z-fibre firings then just as there could not be a world where one had Z-fibre firings and no Z-fibre firings – $\Box \sim(P \,\&t \sim P)$[10] – then, on the assumption of identity, there could not be a world with pains and no Z-fibre firings. Or to put the thought another way, if consciousness were just a physical/functional *concept* then there would be a conceptual bar to thoughts about zombies.

The reason that this could be rejected as an option is not merely that we do we not know that the identity thesis is true yet, but because it seems intuitively right to say that no matter how much evidence that we might accrue about the brain, no matter if we had a complete physics in place, there would still be no *a priori* warrant for ruling out the separability of mental and physical states. And if we cannot rule out a claim *a priori* then it follows that there is no conceptual problem holding it to be the case. I think that there is much more that needs to be said about this – for my own part I suspect that there could be some deep connection between physical states of affairs and phenomenal states that might, on further discoveries, rule out *a priori* the zombie claim. For now, however, we will allow that zombies are conceivable. The question remains then whether they are, accordingly, possible, and I think that the Arnauld objection does at least suggest that a gap might be opened up between conceivability and possibility. Can that gap be plugged?

IV Ideal Conception

Can we refine C➛P so that it delivers infallibly? Some philosophers have appealed to ideal forms of conceivability as illuminating genuine possibilities. David Chalmers,[11] for instance, distinguishes *prima facie* from ideal conceivability, where someone conceiving in the former sense wouldn't have sufficient warrant for the claim that such and such was possible, but someone who conceived of things in the latter sense would. So in Arnauld's confused geometer we clearly had a case of *prima facie* conceivability. Any ideal conceiver would not have made the elementary error of thinking that the triangle conceived of could lack the salient property. Ideal conceivers need, at least, to be experts in the areas to which the conceivings apply. However, this looks just a bit too like Descartes' appeal to clarity and distinctness in our concepts, and wouldn't deal with the following case, which Chalmers mentions. Frege thought that there was a set of all sets, presumably having thought carefully about the matter. Later, however, Russell came along and showed that there couldn't be such a thing because the very idea of a set of all sets generates a paradox. Now Frege, unlike the confused geometer, was an expert in the field in which he had his conceivings.

One might argue in fact that, at the time Frege thought that there was such an object, hardly anyone was better placed to adjudge whether there could be a set of all sets. But, after Russell, we know that such a thing is an impossible object. Here is a case of conceivability failing to reveal possibility, but not, it seems right to say, mere *prima facie* conceivability. This was, however, a case of *secunda facie* conceivability, according to Chalmers, which is, as its name suggests, a step up from our first, ill-considered, conceivings, but falls short of ideal conceivability. Russell, however, was in the position of the ideal conceiver. So what, exactly does it mean to conceive of something ideally? The following might be a first shot at framing such a principle.

> ($C_{irr}\rightarrow$P) If S is conceivable on ideal rational reflection then S is possible.

But what is meant by *ideal rational reflection?* The thought seems to be that a statement S would be ideally conceivable 'if an ideal reasoner would find that it passed the relevant tests',[12] i.e., tests such as attempting to rule S out *a priori*. Chalmers himself doesn't find this entirely satisfactory because of the difficulties faced in making sense of the notion of an ideal reasoner, that is, whether for any imagined 'ideal' reasoner we could imagine one who is even smarter (more ideal), so he instead he appeals to the notion of 'undefeatability by better reasoning.' So we get something like.

> ($C_{\sim rd}\rightarrow$P) If the justification for S cannot be rationally defeated then S is possible.

On this version we can see clearly that Frege lacked warrant for his belief in the set of all sets because better reasoning (Russell's) would have defeated the belief. But this again begs the question when are we in a position to say that we have everything that we need to satisfy $C_{\sim rd}\rightarrow$P? I just can't see how this looks in better shape than $C_{cd}\rightarrow$P.

We believe things when we don't have before us defeaters for our beliefs, so in a sense we get this for free. I think that such and such is conceivable until I am given grounds for not taking it to be so (someone introduces a defeater such as a clear counterexample to my C➛P claim). What we have to work much harder for (and I am pessimistic about reaching this goal) is a justification that we are in the position where no further defeaters are possibly forthcoming - and this just reiterates the question asked all along, 'when are our conceivings infallible guides to possibility?'

So if it looks like there is always going to be a gap between what we conceive and what is possibly the case then does this give us reason to suppose that even were zombies conceivable they needn't be possible? I think so. And if there is always doubt that conceivability hooks onto real possibility then at the very least we can assert that the case against materialism has not been conclusively established on the basis of zombie conceivability alone.

Paul Sperring

Notes

1 Obviously there could be functioning brain states that occurred without any conscious states at all. There are no conscious states that accompany the brain states that are involved in the regulation of my breathing, for instance. I have in mind just the salient brain states that are accompanied, in non-zombie cases, by phenomenal states (in the literature, C-fibres firing accompanied by feelings of pain).

2 By 'Z-fibres' I mean just whatever particular brain states are identified, by the neuroscientist, with the phenomenal experiences had by the possessor of those brain states.

3 The stoutest defence of the argument for dualism from zombie conceivability can be found in David Chalmers' book *The Conscious Mind* (New York: Oxford University Press, 1996).

4 Which isn't to say that what is possible is just defined as what we think about coherently. Possible states of affairs exist independently of our capacity to conceive of them, which Descartes illustrates by appealing to, for instance, complex geometric properties. These properties exist whether I choose to think of them or not, and I am not free not to think about them other than as they are once they are discovered to be essential to the object in question. Possibility, for Descartes, is fixed by metaphysical not epistemological limits.

5 *The Philosophical Writings of Descartes, Volume II* (Cambridge: Cambridge University Press, 1984), p.54

6 It might be objected that if the argument hinges on the truth of theism then it has little chance of going through. However God is not necessary for the argument, Descartes is merely employing God as a device to delimit the space of possibilities (much in the way that modern arguments employ possible worlds talk). So if a certain state of affairs is conceivable then (by whatever power) that state of affairs is possibly instantiable exactly as conceived.

7 *The Philosophical Writings of Descartes, Volume II*, pp.139-143

8 Or one might say, that it is not *apparent* cases of conceivability that entail possibility but *genuine* cases, and Descartes clarity and distinctness criterion provides us with a means of distinguishing between the two. This claim is not without its problems as we shall soon see.

9 There is something in this – we do seem to have unimpeachable warrant for the claim that our existence is indubitable – but whether other claims can be known with quite the clarity and distinctness of the *cogito* is doubtful.

10 Meaning 'necessarily, nothing can be both itself and not itself.'

11 David Chalmers, 'Does Conceivability Entail Possibility?', in Tamar Szabo Gendler and John Hawthorne (eds.), *Conceivability and Possibility* (New York: Oxford University Press, 2002), 145-200.

12 ibid, p.148

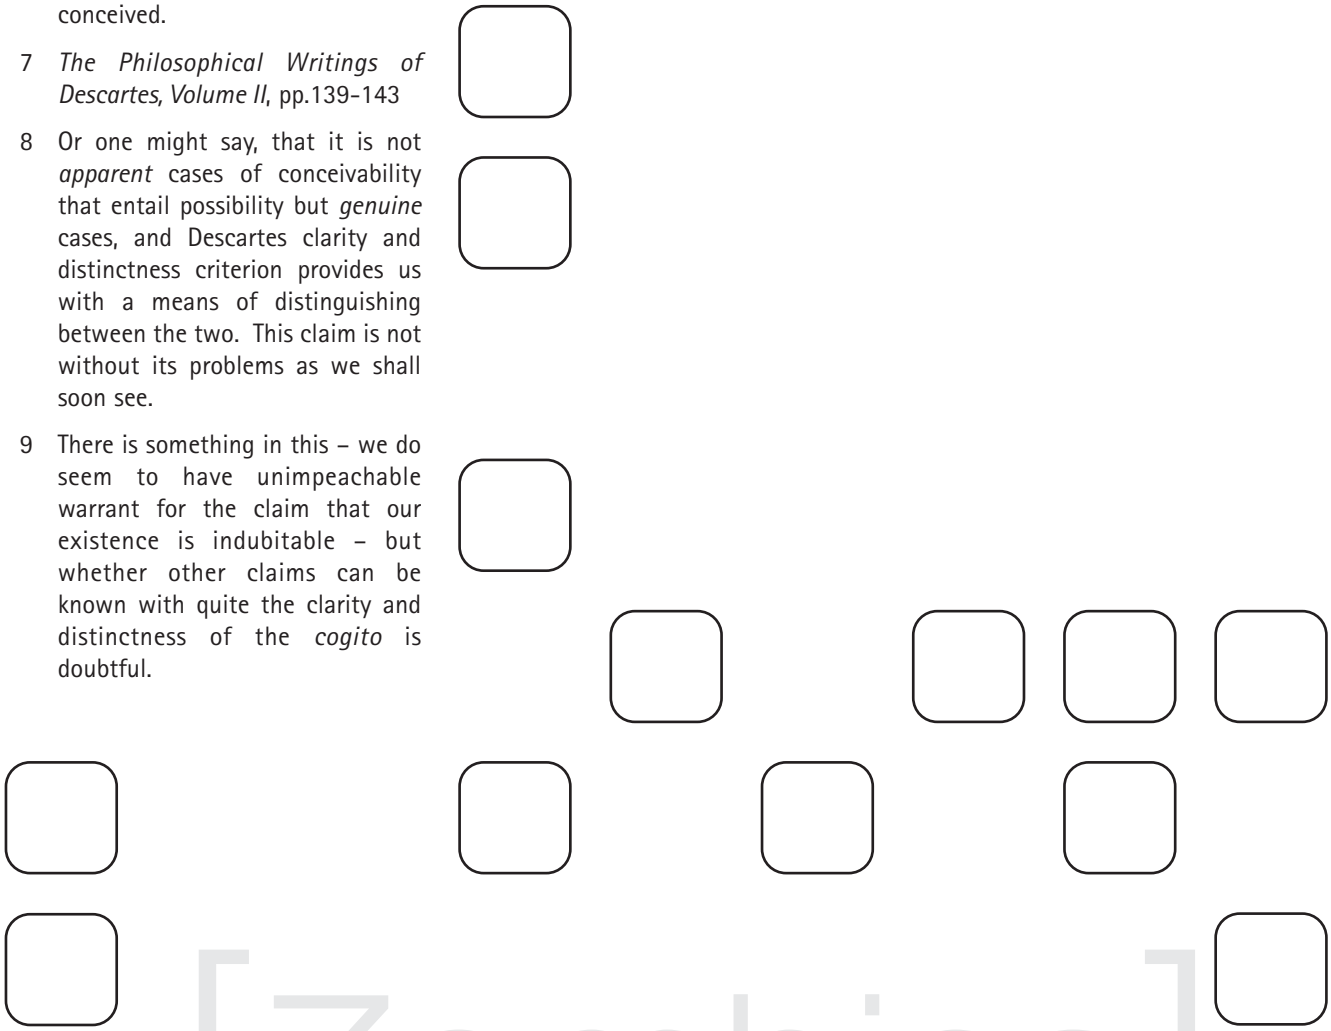

Kathy Behrendt

Despair, Liberation & Everyday Life: Two Bundle Views of Personal Identity

Philosophy sometimes has the reputation of dealing with matters outside the realm of 'everyday life', and trading in ideas that float free from anything beyond the armchair in which we sit contemplating them. I want to talk about a standard armchair-branch of philosophy – personal identity theory – and the real-life effects it either has had or has apparently failed to have upon two philosophers. The two philosophers are David Hume and Derek Parfit. Both arrive at similar and quite radical beliefs about personal identity. And both have documented the difficulty of sustaining these beliefs in their day-to-day lives. For those considering embarking upon philosophical study – whether formally or not – this last point may seem discouraging, reinforcing a picture of a discipline that even on the admission of its own practitioners has little impact on everyday life or concerns. I will explore these two philosophers' views on personal identity in some detail, and outline the conflicts which they claim to exist between their philosophical and non-philosophical thinking. I will go on to propose that these conflicts do not in fact reinforce an opposition between everyday life and philosophy.

David Hume's brief discussion of personal identity has been highly influential.[1] A key aspect of his account is captured by his following, famous remark: 'For my part, when I enter most intimately into what I call *myself*, I always stumble on some particular perception or other, of heat or cold, light or shade, love or hatred, pain or pleasure. I can never catch *myself* at any time without a perception, and never can observe anything but the perception'.[2] Having gone to look for a self in the sense in which it was commonly understood in his time – a single, enduring entity that stays the same throughout our lives – all Hume finds instead is a collection or 'bundle' of constantly and rapidly changing thoughts, sensations, and perceptions. He concludes that, as far as he can see, persons are nothing but bundles of changing mental properties, with no identity over time.

But Hume is very concerned with why we have the ideas that we do, and he cannot deny that we *think* we are enduring persons that survive and persist throughout the many changes that go on in our minds. This belief in personal identity arises, he speculates, from our observation that the various, distinct impressions in my mind succeed one another over time, and resemble one another, and appear to be causally linked – one thought giving rise to another. The result is still a bundle of diverse parts, but a bundle that is tied together by these relations of closeness in time (or as he calls it, 'contiguity'), resemblance, and causation, so as to give rise to the illusion of the identity of ourselves over time. Although philosophical reflection about the matter forces us to admit that there's still no single thing to be found that endures throughout the related impressions, Hume says that 'we cannot long sustain our philosophy'. We will invariably yield to the non-philosophical belief that we are the same thing over time – a self, or soul, or substance, all of which are fictions of our imagination.[3]

This non-philosophical conviction about the identity of the self is held by those who Hume sometimes, rather unflatteringly, calls 'the vulgar'. It is important to understand that the vulgar are not a separate and perhaps lower class of people. Vulgarity is, rather, a state of mind in which we all participate at some time or another, and no one is more ready to admit to vulgar tendencies than Hume himself. And his vulgar tendencies persist in spite of the fact that he clearly recognises that there is no room for his philosophical views within the vulgar frame of mind. This leads him grimly to contemplate the possibility that all our 'reflections very refin'd and metaphysical have little or no influence upon us'.[4] In some of the most personal passages to be found in the philosophical literature of his time, Hume goes on to confess that the apparent contradictions that he has

unveiled between our vulgar fictions (which include our idea of our identity) and his metaphysical reflections (which reveal that the idea is ill-founded) have plunged him into 'the most deplorable condition imaginable, inviron'd with the deepest darkness'.[5] Happily, he notes, nature 'cures me of this philosophical melancholy and delirium' – he describes how he has dinner, gets together with friends, and plays a few games of backgammon, upon which his philosophical quandaries are forgotten, leaving him free to indulge his vulgar attitudes without any feeling of turmoil.[6]

Derek Parfit, writing some 240 years after Hume, came to some similar conclusions about the self.[7] Like Hume, Parfit denies that there is any evidence for a single, enduring self or person. A person is a collection of fleeting, disparate parts. Unlike Hume, Parfit does not restrict his attention to the mind but also includes the brain and the rest of the body amongst those parts. The result is still a bundle theory, since the brain and body are also made up of many different and changing parts. While the molecules of my brain and the rest of my body may not replace themselves at the rate at which my thoughts do, they nevertheless do, all of, them, change or disappear throughout my lifetime, given I live long enough.

In contrast to Hume, Parfit opts for a somewhat more relaxed notion of identity by which, if all the changes in my mind and body continue on as normal, with no radical breaks or divisions (e.g. I don't lapse into a coma, or develop severe amnesia, or have my still-functioning brain divided in half by surgeons and transplanted into the empty skulls of two waiting recipients – a favourite fantasy in personal identity theory), I can be said to remain the same person throughout my life.

What is striking are not Parfit's views about normal personal identity, but certain other conclusions he draws from the Humean starting point that persons come in bundles. These conclusions concern what matters to us with respect to our future existence. 'What matters', in this case, carries a rational and moral sense, meaning that it concerns how we should act towards ourselves and others as moral beings with certain desires that we wish fulfilled. For Parfit, how we should act should be dictated in large part by the metaphysical facts about personal identity – by what kind of things we are. If we recognise that there is for each of us no single thing (like a soul) that endures throughout our lives and which truly constitutes who we are, then we cannot, rationally speaking, be concerned for or worried about the preservation and flourishing of this thing, self, soul, or whatever we call it.

The only reasonable candidates for what ought to matter to us are the ones already mentioned – the brain, the rest of the body, and the mind or psychological content (beliefs, wishes, hopes, aspirations, personality traits, and so forth). Amongst those candidates, Parfit holds that it is the psychological content that is generally the focus of our main concerns about our future; with the possible exception of, say, super-models and athletes, most of us place more importance on our mind continuing on unimpaired in the future than our body. In particular, we want our desires and ambitions fulfilled or at the very least to have the opportunity to fulfil them. And we want our thoughts, beliefs, memories, and personality to continue on – perhaps not in the same way, as we may wish for advancement or improvement, but without any radical breaks or lapses. In short, to use Parfit's terms, we want the future to involve psychological continuity and connectedness to ourselves in the present.

However, Parfit argues, all these things could continue, while we ourselves ceased to exist, or our continued existence became an indeterminate matter. Such a thing may never actually happen, but it is nevertheless a logical possibility. Consider a variation on the previously-mentioned brain division and transplant thought-experiment. Imagine you will divide overnight like an amoeba, into two new, physically-alike persons both of whom will 'inherit' all your psychology – your thoughts, memories, personality and so forth. I say two new persons because any other description seems implausible; we can't say that you remain as one person with two bodies, since both embodied 'halves' of you would be spatially distinct and could go on to

[Identity]

lead disparate lives. They may meet later and have a falling out and one kill the other in a fight – a situation which would seem to make it impossible to talk reasonably of there being one person. Perhaps, you could argue instead, there is only one new person, and you survive as one of the two products of the division, the other half being someone else (a copy perhaps?) But on what basis could it be said that one post-division person rather than the other is really you? They are by definition exactly alike in terms of the features they have at the moment of division. Maybe we could just pick one. Yes, certainly, but this would be an arbitrary decision made for the sake of convenience and not on the basis of any metaphysical fact – metaphysically speaking, as said, the two halves have all the same features.

While we might intuitively be appalled at the prospect of our division, we shouldn't be, according to Parfit. For although one cannot reasonably be said to survive such a division, what matters to most of us does. Our beliefs, hopes, memories, ambitions, and personality, all survive. In this case, they survive twice over. This is nothing like death, it seems, and may even be better than normal life – if you knew this were going to happen to you soon, you could plan to go to university straight away after school and also have a gap year.

One's response to all this may be, so what? Since we don't divide like amoebas the above sort of speculations could not result in new metaphysical beliefs about our actual selves. But this is not so. The amoeba thought-experiment is an exercise to get us to recognise certain facts about ourselves. In normal life, identity goes together with psychological continuity and connectedness, and so we are not forced to choose between them, and

can easily mistake our concern about one for concern about the other. But thought experiments like the amoeba one are meant to show that psychological continuity and connectedness can come apart from identity and exist without it, and so are different from it. Once we recognise this then even in normal cases involving unproblematic identity of ourselves over time, we will, Parfit hopes, realise that it is not this identity itself but the separable psychological continuity and connectedness that matters to us.

These conclusions that Parfit arrives at concerning what matters to us in survival all result from his initial, Humean, metaphysical view that a person is nothing more than a complex entity made up of fleeting and changing parts. If the view that he and Hume rejected were true – if a person were ultimately a single, enduring entity that continued to exist unvarying throughout all these changes – then Parfit could not draw the conclusions that he does. Because such an entity has no parts and so by definition is not divisible, Parfit would

have to conclude that it would be destroyed upon division, or would go with one of the products of the division (we might not be able to tell which one, but there would still be a fact of the matter). But Parfit does not think there are such entities, and so he must accept what he believes are the consequences of the metaphysical view he does hold, and those consequences are that psychological continuity and connectedness matter and personal identity does not.

These consequences are positive, as far as Parfit is concerned. The recognition that there is no single, unchanging self that endures throughout our lives, means that, even in cases of normal lives, what we call our identity over time is a matter of degree. Even if we don't divide or undergo sudden radical physical or psychological change, we do change over time and therefore our present self may be more or less connected with our past or future self, depending upon the extent of changes between them. Realising this can help us to alleviate negative emotions and attitudes such as fear of death, egoism, and regret about the passing of our life, all of which gain their strength, according to Parfit, from the false belief that it is I, the *exact* same person, who is involved throughout. Parfit also sees positive ethical consequences of his view.

My concern about my distant-future self, with whom my psychological connections are weak, will be more like my concern for other people to whom I am not currently connected. Therefore I have reason to act in a more altruistic, less self-interested way; my actions on behalf of my future self should not be radically different than my actions towards other people. While none of these consequences that Parfit discusses are without controversy, the point remains that he at least finds them extremely positive and even, in his word, 'liberating'.[8]

This is in sharp contrast to the feelings of darkness and despair into which we saw Hume driven by his own philosophical views. Even if Hume's dining and gaming tendencies show that his life was not ruled by despair, an appendix he added to the *Treatise* makes matters look even more bleak.

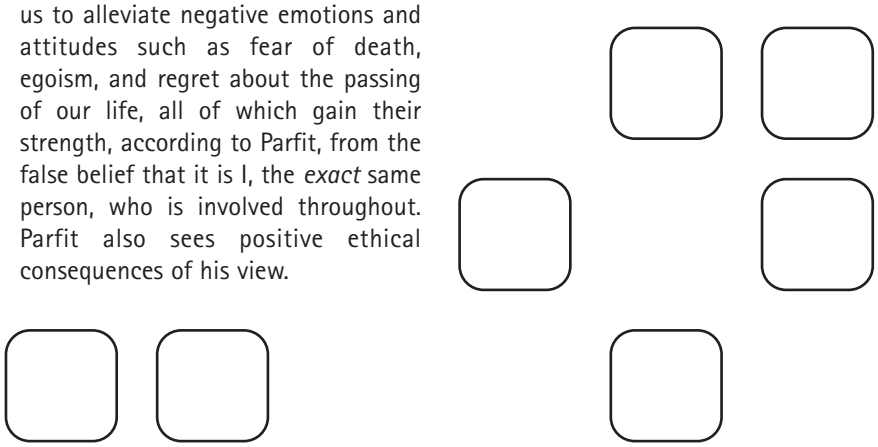

There he makes clear that the relations which bind together the bundle of thoughts and perceptions that constitute a self are not *real*; each of the perceptions are distinct existences that have the potential to exist apart from one another, and human understanding is incapable of discovering real connections between distinct existences. So the bundle falls apart. We have lost both the single, simple soul and the changing but unified bundle of perceptions and, it would seem, have thereby well and truly lost 'our *selves*' in any meaningful sense of the term. Surely this *is* cause for despair. How could anyone who believes it continue on the same as he or she did prior to such a realisation? And yet continue on Hume does. He leads a long and, if his own accounts are to be believed, quite happy life, filled with travel and friendship and the pleasures of the everyday. Perhaps the most obvious explanation for this, is that he simply cannot bear the truth that philosophy has revealed to him, and shuts it out of his mind in order to get on with his life. As we shall see, though, this is not his own explanation.

If the negativity of Hume's philosophical conclusions were the best explanation for his failure to sustain them, then the opposite should hold true for Parfit. Unlike Hume, he discovered a positive outcome could emerge from the same metaphysical starting-point.[9] Inspired by his liberating views about personal identity, Parfit, we might expect, has no reason to attempt to distract himself from them through dining, backgammon or any other means. And yet, very much in keeping with Hume, Parfit also speaks of the considerable difficulty he has in sustaining his philosophical convictions. In passages reminiscent of Hume's contrast

between philosophy and vulgarity, Parfit talks of a split between the 'reflective or intellectual level' of thinking and some 'lower level'.[10] The reflective level concerns the metaphysical truth of the matter and the lower level what we are generally inclined to believe. Parfit confesses that when he forgets or fails to concentrate on his philosophical views he slips back into the sort of egoistic self-concern that is not rational according to his own arguments.[11] But why is this so, if unlike Hume, he finds his conclusions to be positive? Why is even a reputedly beneficial truth hard to believe?

To answer this, it may help to return to Hume, and note that it is not because his philosophical beliefs are depressing that he has trouble sustaining them. Rather he indicates that philosophical beliefs and thoughts – pleasant or unpleasant – are difficult to sustain, producing as they do only 'mild and moderate sentiments'.[12] Parfit at times seems to be of similar opinion; it is not that the truth is good or bad, but that it is hard to sustain any view that conflicts with the status-quo of everyday thinking (or lack thereof) about the matter; our everyday thinking is, as it were, our 'default position', and according to that mode of thinking, each of us is a single, enduring person whose identity is fixed and stable, and not a fiction or a matter of degree (witness your own, likely less-than-cheerful response to the prospect of your division). Hume and Parfit's agreement on this matter is reinforced when we look at the explanations that they both give for our everyday thinking. Hume says it arises from our 'nature' or instinct; like many other things we perhaps ought not, philosophically speaking, to believe, we cannot help but doing so. We are part of nature and as such are subject to its dictates. Parfit, like Hume, believes that our tendency to cling to this notion has something to do with the kind of beings we are. He holds that we human beings are, for better or for worse, in possession of the concept of a single, enduring and unchanging self; it is built into the structure of our thought and so this concept continues to work its influence upon us even once we know better.[13]

To sum up, we have in Hume and Parfit examples of two philosophers who, while passionate and articulate about their philosophical beliefs, nevertheless and upon their own admission cannot easily sustain those beliefs. We also have some indication as to why this is so, namely that it is built into our nature or our conceptual scheme that it could not be otherwise. Does all of this serve to reinforce the sort of view mentioned at the outset, namely that certain philosophical ideas – be they negative or positive – have little to no effect on our everyday lives, but occur, if at all, in a vacuum?

Before we proclaim a complete victory for vulgarity over philosophy, we should be very careful about how we depict their opposition. Hume himself is very careful in this regard, in that although he treats vulgar impulses as 'natural' he does not view philosophy as somehow unnatural. Indeed he remarks that we (many of us at least) are also naturally inclined to philosophise ("tis almost impossible for the mind of man to rest, like those of beasts, in that narrow circle of objects, which are the subject of daily conversation and action'[14]), and so long as we are compelled to do so we ought to do so, and in doing so the thoughts that ensue, though no doubt requiring more mental effort than backgammon, will not be unnaturally forced. The worst thing one can do, as

far as Hume is concerned on this matter, is to have an inclination to philosophise but suppress or reject it. Hume is not simply driven away from philosophy and into the real world – he is catapulted back and forth between them in a relentless though often exhilarating fashion, and by the same natural forces.[15] If we treat philosophical speculation as not the enemy of natural inclination but an extension of it, we can view this swinging back and forth as two interdependent rather than opposed activities. They are interdependent because the loss of one is a loss for the other: clearly, unrelieved philosophical contemplation is not possible and would be detrimental to one's overall well-being as far as Hume is concerned; but likewise the complete denial of philosophical impulses is also damaging. As Hume proclaims, regarding his philosophical impulses, 'these sentiments spring up naturally in my present disposition; and shou'd I endeavour to banish them, by attaching myself to any other business or diversion, I *feel* I shou'd be a loser in point of pleasure; and this is the origin of my philosophy'.[16]

Similarly for Parfit, the conflict between aspects of his philosophical and his non-philosophical thinking

does not constitute a complete separation of them. While Parfit acknowledges the difficulty of sustaining his carefully considered views on personal identity, it in the end amounts to no more than that: a difficulty, but one which he strives to overcome in the interest of the truth. 'The truth is very different from what we are inclined to believe', he is fond of stating.[17] It is therefore harder won. It is perhaps not surprising then that we often opt for the easiest way of getting by, neglecting or setting aside beliefs which are apparently true. But the fact that such beliefs are harder won and as a consequence perhaps more rare is not reason to view them as not properly belonging to and impacting on our lives as a whole. Parfit makes clear that despite his lapses, his views on personal identity have had a profound affect on his life: prior to establishing those views, he writes, 'I seemed imprisoned in myself. My life seemed like a glass tunnel, through which I was moving faster every year, and at the end of which there was darkness. When I changed my view, the walls of my glass tunnel disappeared. I now live in the open air.' [18]

I started out with a contrast between philosophy and everyday life. But I think we should now see this as a false opposition. It is without doubt that any view which suggests that our identity is a fiction, or is something that does not matter, is extremely daunting. It is possible that even if we concentrated all our effort upon sustaining such a view we would on occasion fail.[19] This failure may, as speculated, have less to do with the nature of the view in question than with our human nature or the way our minds are constructed. Whatever the reason, the fact that both Hume and Parfit find it difficult and sometimes impossible to sustain their philosophical convictions or act purely in accordance with them does not detract from the fact that these convictions are a fundamental part of their whole lives, which would be poorer without them.

Notes

1 The discussion can be found in Hume, *A Treatise of Human Nature*, D.F. Norton and M.J. Norton, eds. (Oxford: Oxford University Press, 2000), Book I, Part 4, Sect. 6. There are also some brief but significant remarks on the subject to be found in his Appendix to the *Treatise*.

2 Hume, pp. 164-165.

3 Hume, p. 166.

4 Hume, p. 175. The whole of the Conclusion of Book I of the *Treatise*, from which this and the following passages are taken, is an eloquent examination of the place of philosophy in life.

5 Hume, P. 175.

6 Hume, p. 175.

7 Parfit's most detailed and best-known account of his views on personal identity is to be found in his *Reasons and Persons* (Oxford: Oxford University Press, 1984), Part Three, and all my citations of Parfit are from that work, unless mentioned otherwise.

8 Parfit, p. 281.

9 See Parfit p. 282 for his view on the difference between Hume's conclusions and his own.

10 Parfit, p. 279.

11 Parfit, pp. 279-280.

12 Hume, p. 176.

13 This is a point which he develops in 'The Unimportance of Identity', in H. Harris (ed.) *Identity* (Oxford: Clarendon Press, 1995), pp. 13-45.

14 Hume, p. 176.

15 Again, see his Conclusion of Book I of the *Treatise* for a detailed account.

16 Hume, p. 176.

17 Parfit, p. 281.

18 Parfit, p. 281.

19 Though Parfit suggests that Buddhism may be the nearest thing to a full-time implementation of his view; see *Reasons and Persons*, pp. 273, 280, and Appendix J.

Christopher Norris

[Anti-Realism]

Scepticism and the Limits of Sense

I

Michael Dummett has been occupied over the past four decades in exploring, refining, and (mostly) defending an anti-realist approach to various fields of knowledge or branches of enquiry.[1] Anti-realism, on Dummett's account, is defined chiefly in negative terms, i.e., by its denial of certain theses that he takes to characterise the realist position. For the realist there is a large class of statements whose truth-value is strictly undecidable since it lies beyond our utmost powers of verification or falsification yet concerning which we can rightfully assert that they must be *either* true *or* false – objectively so – despite our lack of knowledge concerning them. What decides that value is the way things stand in reality, that is, the existence of certain *truth-makers* (facts, circumstances, real-world [including historical] events, mathematical or other such abstract verities) to which those statements correspond in their role as *truth-bearers*. Truth is conceived as recognition-transcendent in the sense that it depends not at all on the scope and limits of our cognitive or epistemic powers. For the anti-realist, conversely, any truth-apt statement has to meet the condition that its truth-value can be specified in terms of some available proof-procedure or method of verification. To suppose otherwise is to believe – nonsensically – that we could somehow acquire or manifest a grasp of what it takes for that statement to be true (or false) while lacking just the kind of knowledge required to decide the issue either way. In which case we should think of truth as 'epistemically constrained', or of statements as possessing a truth-value only in so far as we can (or at any rate could in principle) find it out by some investigative means. The realist must therefore be deluded – metaphysically out on a limb – if he or she asserts the existence of truths that would lie beyond our utmost cognitive, epistemic, or probative reach.

Dummett's other chief claim to originality is to have clarified this whole debate by posing it in logico-linguistic terms or by placing it on ground that has been worked over most thoroughly by philosophers of logic and language in the post-Fregean line of descent. Thus, as he wrote in 1978, '[t]he whole point of my approach . . . has been to show that the theory of meaning underlies metaphysics. If I have made any worthwhile contribution to philosophy, I think it must lie in having raised the issue in these terms'.[2] And again, in a retrospective piece some fifteen years later on:

> [t]he opinion is sometimes expressed that I succeeded in opening up a genuine philosophical problem, or range of problems, but that the resulting topic has little to do with traditional disputes concerning realism. That was certainly not my intention: I meant to apply a new technique to such wholly traditional questions as realism about the external world and about the mental, questions which I continue to believe I characterised correctly.[3]

'Correctly' is somewhat ambiguous here as between 'getting the issue into a more perspicuous focus without any bias either way' (Dummett's professedly neutral or even-handed line of approach) and 'presenting that issue so as to highlight the problems with realism' (which is how that approach most often works out in particular contexts of debate). For the regular upshot of Dummett's analyses is to cast the realist as defender of an over-committed metaphysical doctrine and hence to treat anti-realism as the default option for anyone who would wisely seek to shuck off such excess philosophic baggage. Where the realist errs is in supposing that we could ever conceive the existence of truths that surpassed our best powers of ascertainment. This follows – so he argues – from certain crucial considerations about the operative scope and limits of human understanding as embodied in our various, linguistically articulated

means of acquiring and manifesting such truths. What thus becomes plain is the sheer impossibility that our truth-predicates might have some valid application to statements for which we lack any adequate proof-procedure or means of verification, yet whose well-formedness leads us to think that they *must* be either true or false – objectively so – quite apart from such issues of epistemic warrant.

This applies just as much to logic, mathematics, and the formal sciences as to areas of investigation (such as physics or history) where the relevant constraints are chiefly those of empirical or evidential warrant. Thus for instance, as regards mathematics, Dummett adopts an intuitionist approach according to which provability (not objective truth) is the sole criterion and we are therefore wrong to claim of any well-formed yet so far unproven theorem or conjecture that it *must* be either true or false despite its undecidability by the best means at our disposal.[4] To this extent Dummett follows Frege and the later Wittgenstein – albeit with certain express reservations – in arriving at his anti-realist position on issues in the philosophy of language and logic. What he takes from Wittgenstein is a generalisation of Frege's 'context principle', that is, the idea that terms can only have meaning in the context of some given proposition, and hence – extending this principle – that the meaning of that same proposition can itself be construed only with reference to the conditions of verifiability which apply to propositions of just that type within a certain area of discourse.[5] There is a tension in Dummett's argument here since he rejects any radically holistic or contextualist theory of meaning on the grounds that it cannot explain how we could ever acquire or manifest a grasp of this or that particular proposition, as would seem prerequisite for our coming to understand its role within any such wider context. Thus Dummett declares very firmly in favour of a logico-semantic approach based on the principle of compositionality, i.e., the principle that sentence-meaning can be specified in terms of those various component parts (subjects, predicates, logical connectives, etc.) that between them serve to identify its sense and reference. All the same Dummett's anti-realism can be seen to push a long way in that other, more extreme contextualist direction since it entails the idea that statements can be taken as meaningful or truth-apt (more precisely: as candidates for 'warranted assertibility') only on condition that they play some role in our shared practices or accepted methods of proof and verification. On this view – to repeat – we could never be justified in asserting with regard to some particular statement that it *must* be either true or false as a matter of objective (i.e., verification-transcendent) fact even though we lack the evidential means to ascertain its truth-value.

For if indeed it is the case, as Dummett argues, that assertoric warrant extends just so far as the range of statements for which we possess – or might come to possess – decisive evidence either way, then objectivist talk of truth or falsehood is simply off-bounds for statements of the so-called 'disputed class', i.e., those that are undecidable to the best of our knowledge. Rather such statements are neither-true-nor-false since they exceed the scope of warranted assertibility as defined by criteria which cannot but be those of shared understanding – whether within some relatively wide or relatively specialised community – with regard to what should properly count as an instance of proof or verification.

Dummett has two chief arguments to this effect, both of them taken (by himself and others) as central to the anti-realist case. The 'acquisition argument' maintains (after Wittgenstein) that warranted assertibility is a matter of our learning to apply the relevant criteria within this or that linguistic-communicative context, while the 'manifestation argument' further requires – again after Wittgenstein – that we show ourselves competent to exercise that grasp by engaging in various communally recognised forms of behaviour, expression, or rule-governed practice.

On both counts, therefore, it cannot make sense to posit the existence of truths that lie beyond our capacity to produce evidence for them, evidence which qualifies as such according to the norms of assertoric warrant that define the scope and limits of attainable knowledge. At this point one should perhaps acknowledge that Dummett sets out on his own submission not so much to argue the case for anti-realism as to test its applicability – along with that of the rival (realist) hypothesis – across different areas of discourse. All the same one may reasonably doubt these claims of neutrality or even-handedness when set against Dummett's very evident bias in favour of anti-realism, that is to say, his frank inability to conceive what the realist could possibly *mean* by upholding the existence of objective truth-values for unprovable hypotheses or statements belonging to the disputed class. 'For the anti-realist', he remarks, 'an understanding of [any] statement consists in knowing what counts as adequate evidence for the assertion of the statement, and the truth of the statement can consist only in the existence of such evidence'.[6] From which it follows necessarily – on Dummett's account – that '[t]he notion of truth, when it is introduced, must be explained, in some manner, in terms of our capacity to recognise statements as true, and not in terms of a condition which transcends human capacities'.

Thus in his view it is self-contradictory to claim – as if we could somehow *know* this to be the case - that there exist certain truths for which we lack any means of verification or whose truth-value is beyond the grasp of creatures such as ourselves with our particular range of sensory inputs, perceptual modes, cognitive powers, capacities of formal reasoning, and so forth. In which case statements of the 'disputed class' are exceptions to the logical law of bivalence which holds that they must be *either* true *or* false regardless of whether we are now (or might ever be) so placed as to decide the issue. On the contrary: such statements must be taken not only as neither true nor false to the best of our knowledge but as neither true nor false *sans phrase*. Thus Goldbach's conjecture (that every even number is the sum of two primes) may well have been tested up to huge numerical values on the most powerful computer programmes and may also possess the utmost degree of intuitive conviction but must still – since lacking any formal proof – be counted neither true nor false.[7] Or again, take the case of a speculative astrophysical statement such as: 'There exists a duplicate solar system in some epistemically inaccessible region of the expanding universe' (i.e., too remote and receding too fast for its electro-magnetic signals to reach our terrestrial radio telescopes).[8] Here again, according to Dummett, we shall breach the requirement of warranted assertibility – and lapse into incoherence – if we say: 'Well, the statement is either true or false as a matter of objective fact even though we earthlings will never find out barring some (at present) inconceivable advance in our means of observation'.

This latter example brings out the kinship between Dummett's logico-semantic version of the anti-realist case and the stance adopted by verificationists in epistemology and philosophy of science. On their view we cannot be justified in venturing beyond the best empirical evidence and asserting the existence – the objective reality – of certain items (such as remote galaxies or elusive subatomic particles) whose role in our present-best scientific theories licences at most a non-committal attitude in that regard.[9] This position – first adopted by the great nineteenth-century physicist Ernst Mach with regard to the existence of atoms – has lately received a powerful re-statement under the title 'constructive empiricism' by Bas van Fraassen.[10] Its affinity with Dummett's line of argument comes out very clearly when van Fraassen contrasts his own outlook in matters scientific and philosophical with that of his (presumptively misguided) realist opponent. For the latter, he writes, 'science aims to give us, in its theories, a literally true story of what the world is like; and acceptance of a scientific theory involves the belief that it is true'.[11] For the constructive empiricist, on the other hand, 'science aims to give us theories which are empirically adequate;[12] and acceptance of a theory involves a belief only that it is empirically adequate'. Where van Fraassen most strikingly differs with Dummett is in making no pretence of judicious even-handedness as between these two doctrines and adopting a strong, even (at times) a downright contemptuous attitude toward the former. Thus scientific realism invites the charge of 'empty strutting and posturing', of putting up a false 'display of courage not under fire', and moreover of 'avow[ing] additional resources that cannot feel the pinch of misfortune any earlier'.[13] This is because, as van Fraassen sees it, realism *claims* to 'answer more questions' and to give us a 'richer, fuller picture of the world' while in fact doing no such thing (since based upon just the same range of empirical evidence) and moreover taking no additional risks (since subject to just the same chances of empirical disconfirmation).

Hence the odd tone of prosecuting zeal – even of moral repugnance – that tends to overtake van Fraassen's otherwise equable and good-humoured prose when the realist opposition comes into view. Perhaps it may also be explained in part by the range and force of those various counter-arguments that are marshalled against his position. It is most often challenged in current debate by the advocates of 'convergent realism' and 'inference to the best explanation', both of which claim to mount a strong rebuttal (if not a logical refutation) of the anti-realist case.[14] On their account realism is a theory with its own well-established scientific credentials, and one that can be tested in just the same way that first-order scientific theories are tested, i.e., through its managing or failing to provide the best, most rational explanation of how and why various branches of science have produced such a likewise well-established range of descriptive, predictive, and causal-explanatory hypotheses. All this evidence must count for nothing – so the argument goes – if we follow van Fraassen and adopt a 'strong' constructive-empiricist approach that refuses to credit the existence of entities (whether subatomic particles or light-bending galaxies with massive gravitational fields) beyond our best means of direct, unaided, or technologically unassisted observation. However we shall then be able to adduce no plausible account of how science has typically advanced through the stages of (1) pure speculation with regard to (e.g.) the existence of atoms, (2) theoretically-supported conjectures wherein they acquire a crucial explanatory role, and (3) the advent of new, more powerful or refined technologies whereby they can either be observed or manipulated, as is the case with atoms nowadays.[15] Besides, there is something grossly anthropocentric about van Fraassen's idea that the limits of unaided *human* observation (more precisely: the limits of what we can observe through 'basic' instruments such as optical microscopes and telescopes rather than advanced instruments like electron microscopes and radio telescopes) should somehow decide what properly counts as an item of physical reality.[16] Thus the realist will remark how much more accurate and powerful are these latest technologies; that we understand their workings well enough to make due allowance for any inbuilt distorting or disturbance effects; and – not least – how van Fraassen's appeal to unaided (or 'naked') observation ignores the sheer amount of perceptual and cognitive processing that goes on between the impact of photons on our retina and the experience of visual images.[17]

Also (just to drive the point home) it is a strange theory which obliges its holder to maintain that some remote celestial body may be taken as real just so long as an astronaut could get close up enough to observe it 'directly' through her spacecraft window – or perhaps through a crude optical telescope – while relinquishing that claim (and figuring merely as a product of empirical convenience) if observed from earth by the most sophisticated means at our present disposal. All of which arguments the realist will take as bearing out her case for scientific realism as a matter of inference to the best (most rational) explanation.

Needless to say, the constructive empiricist will remain staunchly unimpressed by such objections, just as the Dummettian anti-realist will see no force to any counter-claim that the existence of objective (recognition-transcendent) truths is a precondition for our grasp of what constitutes knowledge and progress in mathematics, the physical sciences, and other regions of enquiry.

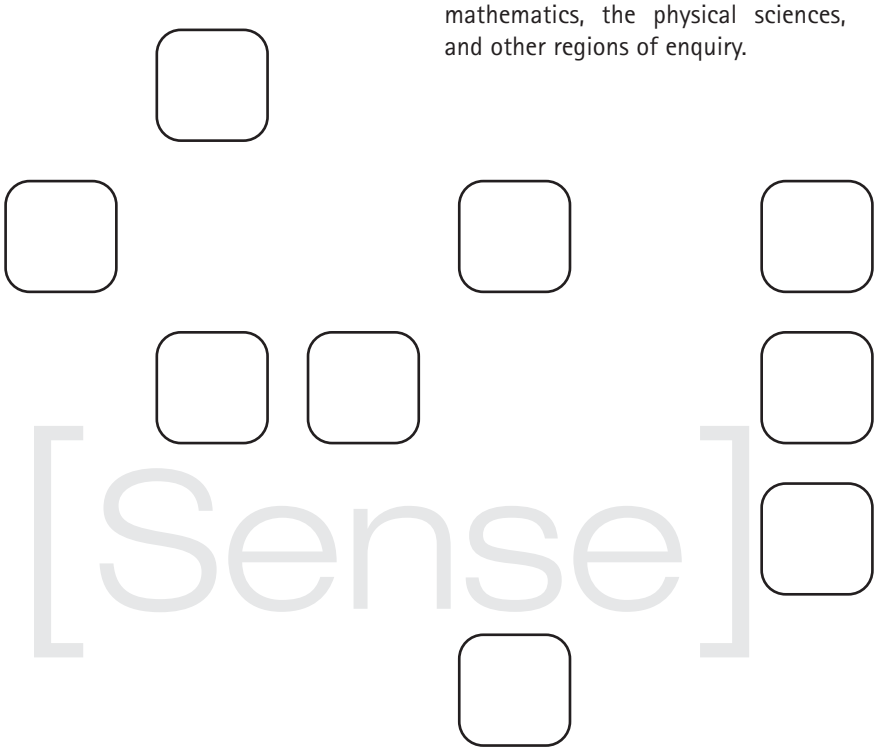

Thus the argument for convergent realism – that terms in a mature scientific theory 'typically refer' and that the laws in such a theory are 'typically approximately true' – will strike the constructive empiricist as a mere fudging of the issue, and besides, as ignoring the sheer range of candidate items (phlogiston, caloric, the luminiferous ether, the planet Vulcan, etc.) which once appeared to meet exactly those requirements but have now passed into the history of discredited scientific lore.[18] To which the convergent realist may respond by pointing out that this 'sceptical meta-induction' (or generalised 'argument from error') plainly fails to work since it presupposes what it sets out to deny, i.e., the fact that our knowledge has advanced to a stage where we can confidently say of such terms – and any putative laws associated with them – that they are empty or non-referring.[19] Also there is the more nuanced version of this argument which distinguishes between totally obsolete theories (like those involving 'phlogiston' or 'the planet Vulcan') and theories which, although strictly false, can be seen to have paved the way for subsequent developments that still hold a place in our current-best scientific thinking. Such would be the case as regards Black's 'caloric' hypothesis since it led on to the theory of specific heat, and likewise as regards the 'luminiferous ether' since – with a somewhat greater stretch of charitable hindsight – we can take it as referring to something very like Maxwell's electro-magnetic field.[20] However, as I have said, these realist rejoinders will cut no ice with the anti-realist or constructive empiricist for whom they will appear just a kind of metaphysical extravagance, that is to say, a needless (and explanatorily vacuous) yielding of hostages to future scientific fortune.

II

I should not wish to give the impression that Dummettian anti-realism and van Fraassen-type constructive empiricism are two variants on the same sceptical theme, or that they don't involve significant differences of argument and emphasis. Dummett's is in one sense a more cautious verificationist approach, arguing its case on primarily linguistic (or logico-semantic) grounds and rejecting – or at any rate purporting to reject – any fixed anti-realist *parti pris* as concerns some particular area of discourse. To this extent it contrasts with van Fraassen's doctrinaire insistence on the folly or the false display of 'courage not under fire' indulged by realists who in truth risk nothing more than straightforward, honest empiricists should their theories at length prove wrong or their putative referents (like 'phlogiston' or 'Vulcan') turn out not to exist. On the other hand there is something of mock humility about Dummett's claim to be merely trying out the rival (realist and anti-realist) hypotheses across a range of areas – from mathematics to morals – with no preconceptions either way. For if taken at anything like full strength (as it often demands to be taken) then Dummett's logico-semantic approach goes much further toward undermining certain basic realist or objectivist conceptions than does van Fraassen's relatively specialised focus on issues in philosophy of science. This difference comes out with particular force when Dummett declares – on precisely such logico-semantic grounds – that any 'gaps in our knowledge' must also be construed as 'gaps in reality', i.e., that if we lack sufficient evidence or a reliable means of verification for some given (e.g., historical) statement then *ex hypothesi* that statement possesses no determinate truth-value and is hence referentially void.

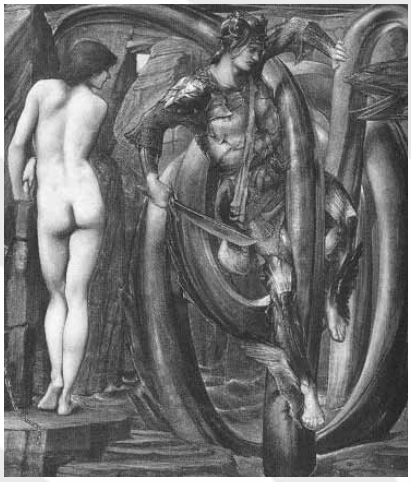

This idea is troublesome for Dummett since he knows very well – as one whose moral and political convictions have led to him to engage actively in campaigns against racist movements like the National Front – that such thinking might fall in with the purposes of right-wing revisionist historiography or even such flagrant abuses as Holocaust-denial.[21] After all, if his argument goes through then it is a fallacy to hold that there are certain claims about the past whose veridical status is a matter of objective (verification-transcendent) truth and which could therefore in no way be affected by any change in our state of knowledge, e.g., by the loss or destruction of evidence or by some large-scale, highly successful programme of ideological brainwashing. The issue is somewhat complicated here by Dummett's frequent suggestion that anti-realism is the best (indeed only) way to keep a grip on such facts since it offers an alternative to the realist's scepticism-inducing idea that truth can always come completely apart from our evidential sources or means of verification. Thus:

> [r]ealism about the past entails that there are numerous true propositions forever in principle unknowable. The effects of a past event may simply dissipate . . . . To the realist, this is just part of the human condition; the anti-realist feels unknowability in principle to be simply intolerable and prefers to view our evidence for and memory of the past as constitutive of it. For him, there cannot be a past fact no evidence of which exists to be discovered, because it is the existence of such evidence that would make it a fact, if it were one.[22]

However this passage shows very clearly that anti-realism, so far from preserving a reliable link between present knowledge and the truth of past events, in fact cuts in just the opposite direction since it renders such 'truth' entirely dependent on various contingent factors including the survival of documentary sources or their having come down to us without suppression or ideological tampering. Thus when the realist takes it as 'just part of the human condition' that 'the effects of a past event may simply dissipate' she is not for one moment suggesting that *past events themselves* – or the truth-value of our statements concerning them – must likewise be thought subject to attrition through factors such as cultural memory-loss or destruction (whether by accident or design) of the relevant information sources. On the contrary: her point is that such statements – including those of the Dummettian 'disputed' (well-formed though undecidable) class – have their truth-value fixed objectively by what *did or did not occur* as a matter of historical fact and quite apart from any gaps, lacunae, or distortions in the documentary record. This places her in sharp opposition to the anti-realist for whom 'unknowability in principle' is felt to be 'simply intolerable' because it leads us to suppose that there may be truths now or forever beyond our epistemic ken.

Hence Dummett's (on the face of it) quite remarkable statement that, to this way of thinking, 'there cannot be a past fact no evidence of which exists to be discovered, because it is the existence of such evidence that would make it a fact, if it were one'.[23] To be sure, there is some room for debate as to just how far this statement goes in a radically anti-realist direction, i.e., toward claiming that the *truth about* – rather than merely our knowledge concerning – past events is a matter of our best available evidence for them. After all, many philosophers nowadays would reject the view – most famously held by Bertrand Russell – that 'facts' are objects (or complexes of objects and properties) which exist 'out there' in the world and which render our statements true or false to the extent that those statements succeed or fail in corresponding to the relevant facts.[24]

Thus it is often remarked – following the widespread 'linguistic turn' whose sources include Frege, late Wittgenstein, and of course Dummett himself – that facts exist only in and through language (i.e., as articulate statements of this or that kind), and hence that any talk of 'correspondence' between statements and facts is at best redundant and at worst downright nonsensical.[25] So one might just construe Dummett as making the more moderate anti-realist, indeed (in a sense) realism-compatible claim that our linguistically articulated *knowledge* of 'the facts' is epistemically constrained or subject to the scope and limits of evidential warrant. Yet this moderate interpretation cannot stand up when set against Dummett's further remark that the anti-realist's refusal to tolerate 'unknowability in principle' must incline him or her 'to view our evidence for and memory of the past as constitutive of it'.

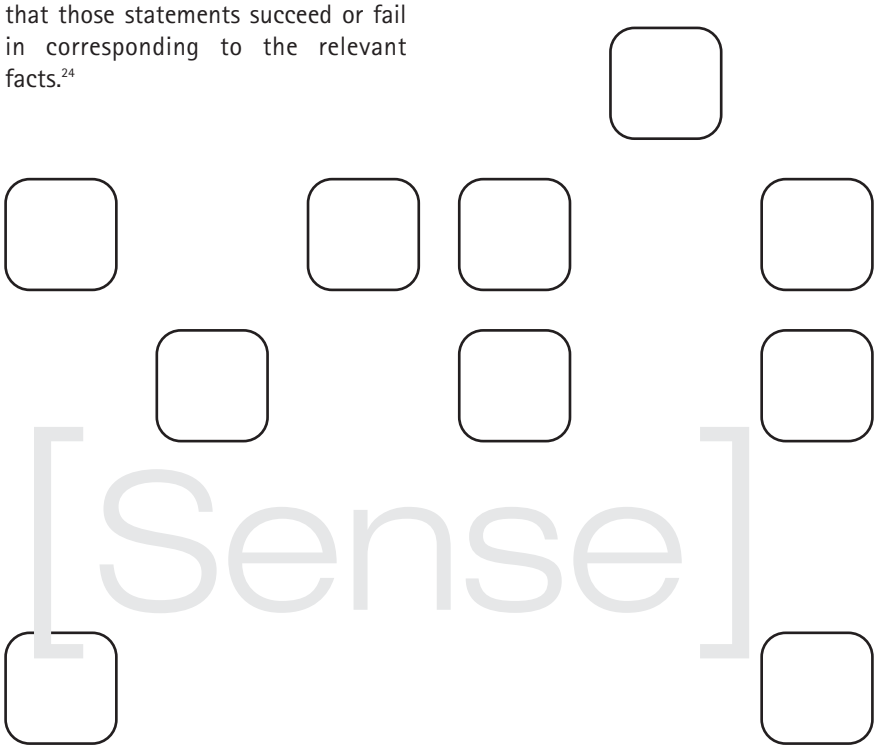

For unless Dummett has carelessly misspoken himself here – omitted to add some crucial qualifying clause – then clearly it is 'the past' (past events themselves rather than our knowledge of them) that should be thought of as somehow *constituted by* whatever evidence lies presently to hand or whatever we are able to retrieve from the data of collective or individual memory. In which case Dummettian anti-realism must be seen to push the linguistic turn to a point where it entails the radical dependence not only of 'the facts' (linguistically conceived) on our state of knowledge concerning them but also of historical truth *per se* on those same 'facts' as recorded, recollected, or evidenced to the best of our ability. And it is then hard to see – on this somewhat disconcerting though textually warranted version of the claim – how Dummett's argument could well stop short of endorsing the idea that present (or future) changes in the nature of our evidence might retroactively affect the occurrence, non-occurrence, or outcome of some past event.

As I have said, Dummett is keenly aware of the affront to all our standing philosophical (as well as everyday-common-sense) convictions represented by this line of thought. Also there are strong counter-arguments – such as that from the existence of 'truth-value links' between past and present – which would seem to give adequate reason for rejecting the idea that any truth of the matter with regard to historical events must be thought of as dependent on our still having access to the same range of evidence as fell within the ken of well-placed observers at the time. These arguments involve the simple device of taking some given statement and supposing it to be spoken at different times with reference back and forth between its differently tensed (but logically equivalent since strictly interchangeable) truth-conditions.[26] Thus, for instance, any statement to the effect 'There was a thunder-storm in Cardiff on April 9th 1987' is true today if and only if 'There is thunder-storm happening right now' was true at some time during April 9th 1987. And likewise, any statement uttered on April 9th 1987 to the effect 'There will a thunder-storm on September 1st 2003' will itself have been true if and only if the statement 'There is a thunder-storm happening right now' is true at some time during September 1st 2003. In which case, it would seem, the anti-realist must be hard put to sustain his thesis in the face of a realist counter-argument which assumes nothing more than the kind of consistency that anyone – whatever their particular views on this question – must surely accept on pain of embracing a straightforward logical absurdity.

Bernard Williams makes a kindred point when he discusses the relationship between myth and history in ancient Greek thought and the way that this relationship can be seen to have changed during the period from Herodotus to Thucydides.[27] What emerged was a new conception of objective time that tended increasingly to separate out these two modalities of discourse and apply more stringent criteria of truth to the various sources – material evidence, documentary (written) reports, first, second or nth-hand oral testimony, folk-memory, 'once-upon-a-time' allusions to a past age of gods and legendary heroes, etc. – which the historian was now called upon to pass in critical review. Williams cites the well-known passage from Thucydides' opening chapter where he impugns the veracity of poets such as Hesiod who conflated mythic with (pseudo-) historical narrative and also of those 'logographers' – Herodotus presumably among them – who failed to draw such distinctions with adequate rigour.[28] This critique carries a strong implication that there is no room within historical discourse, strictly speaking, for the kinds of 'indeterminate' person or event whose existence or occurrence had hitherto occupied a temporally distant twilight zone concerning which chroniclers had felt no need to decide whether (say) Minos, legendary King of Crete, was a god or a human being, and whether his exploits belonged to the realm of a historical myth or demythologised history. What enables this transition is the advent of a new, more objective concept of time whereby people learn to extrapolate from their immediate (intuitive or experiential) grasp of past, present, and future to a longer-term sense of the temporal relations – or the truth-value links – which constitute the historical domain. In Williams' words:

> We become conscious of our being, in temporal terms, some people among others, and with this comes the idea that some of our past was other people's present, that our present was other people's future, and so on; in particular, that what for us, now, is the remote past, for past people was the recent past or the present . . . . [Thus] it has to be recognised that one cannot implicitly treat the remoter past as a peculiar area in which indeterminate happenings and people could exist. If one can say only indeterminate things about them, then that is a matter of our relation to them. Either there was no time at which they existed, so

> they did not exist at all, and are mere stories; or they were as real, and as determinate in their time as similar things are in ours, and we simply do not know enough about them.[29]

I have cited this passage at length partly because – in conjunction with Williams' remarks about Herodotus and Thucydides – it puts historical flesh on the formal argument from truth-value links, and partly because it stands in such sharp contrast to Dummett's understanding of these matters. The anti-realist, we recall, 'feels unknowability in principle to be simply intolerable' and thus prefers 'to view our evidence for and memory of the past as constitutive of it'. For him, moreover, 'there cannot be a past fact no evidence of which exists to be discovered, because it is the existence of such evidence that would make it a fact, if it were one'.[30]

Williams makes no explicit reference to Dummettian anti-realism in this particular context. However one can see that their arguments are opposed point-for-point on all the relevant issues, including what Williams regards as the progress that came about when historians acquired an objective conception of time and – in direct consequence of that – an objectivist (truth-based and critical) conception of their own subject-domain. After all, 'once we accept the idea of historical time, it is quite clear that the gods are essentially indeterminate, in many respects, and could have no fixed or clear relations to it'.[31] In which case there is a sharp distinction to be drawn between such 'indeterminate' (since mythic or temporally unlocated) beings and those 'gaps in reality' which, according to Dummett, result from 'gaps in our knowledge'. Where the latter claim is plausible only in so far as one renounces any notion of objective historical truth the former makes sense only on condition that historical (as opposed to mythic) personages and events be thought of as having existed or occurred quite apart from our evidence or lack of evidence for them.[32] Thus the formal argument from truth-value links can be extended, refined, and filled out in detail so as to offer good reason for doubting the credibility of an anti-realist approach to issues of historical truth. Moreover one could put the case that anti-realism in this current, no matter how sophisticated logico-semantic guise is a reversion to something very like the stage of proto-historical enquiry that Williams locates in the period just before Thucydides developed the methods and techniques of critical historiography. Thus it gives up the idea of objective (verification-transcendent) truth, along with that of a linear, i.e., non-cyclical temporality with truth-value links between past, present, and future. Williams makes this point rather nicely when he remarks that 'Herodotus had also shrewdly discussed the material remains of past times, such as the many wonderful things he saw in Egypt'. However, he continues, 'there is a special, and very typical, twist in Thucydides', when 'assessing the remains of ancient Mycenae that were to be seen in his time, he compares them with the remains that he supposes might be left to future generations by contemporary Athens and Sparta'.[33] What is required for this is the grasp of an objective temporal sequence that stretches back and forward beyond the limits of personal experience yet which takes such experience as its basis for asserting the reality of past events – quite apart from our knowledge concerning them – and the awareness of a future when historians' claims with regard to some presently existing state of affairs will likewise be rendered true or false (whatever their evidential warrant) by the facts of our current situation. In short, 'the explanatory unity of the world binds not just the past and the present, but the present and the future as well; and concrete expression is given to the idea that our today will be someone else's distant past'.[34]

Where anti-realism signally fails to convince is in offering no plausible explanation of how historiography could ever have advanced beyond its stage of confinement to mythic, uncritical, or taken-for-granted modes of communal belief. Indeed, by denying (or finessing) the argument from truth-value links and preferring, as Dummett says, to take 'our evidence for and memory of the past as constitutive of it' anti-realism reverts – in theory at least – to something very like that stage.[35]

Dummett anticipates this objection and goes various ways around in attempting to head off its strong intuitive force. The anti-realist may begin by remarking that it is warranted assertibility, not truth, that is in question here and then go on to argue that realist errs by ignoring the temporally indexed character of what counts as warranted assertibility from one such temporal context to another. That is to say, she (the realist) deploys the apparatus of tense-logic in a merely abstract or formally regimented way without taking sufficient account of the various possible changes, e.g., expansions or contractions in the range of available evidence that may occur with the passage of time. Thus she assumes that the relevant truth-conditions can be specified without substantive or more-than-notional restriction to the particular time of utterance and the kinds of epistemic warrant obtaining at just that time. In which case the anti-realist will demand that their opponent accord a more central role to the agency of time and not assume a static (fundamentally atemporal) conception wherein truth is thought of as evidence-transcendent or epistemically unconstrained. However, as we shall see, this response to the realist's challenge allows Dummett no exit from the paradox of retroactive truth-conferral and indeed involves him in some fairly extravagant conjectures of just that sort.[36] Among them is the idea that in certain (albeit unusual) cases a change in our knowledge of (or evidence for) past events may be thought of as somehow *bringing it about* that those events either should or should not have occurred, or transpired in some particular way.

At this point the realist will most likely reply that if anti-realism lends credence to such patently absurd ideas then they had best be seen as a *reductio ad absurdum* of the anti-realist case, and hence more usefully employed in showing just what's wrong with Dummett's logico-semantic update on verificationist themes. Thus it is no great distance – 'logically' speaking – from the thesis that truth-values cannot possibly transcend the limits of verification or assertoric warrant to the notion that the 'truth' of past events must indeed be subject to (even in some sense determined by) whatever we possess in the way of corroborative evidence for them. Here again Dummett is aware of the obvious realist rejoinder, i.e., that ascriptions of truth differ from ascriptions of empirical warrant, justified belief, present-best knowledge, and so forth, since truth-values are strictly indefeasible by any evidence that might turn up (or drop out) in the course of further enquiry. Still he feels compelled to adopt an anti-realist position – and to accept at least some of those awkward consequences – on logical as well as metaphysical grounds. That is to say, Dummett simply cannot make sense of the basic realist claim that we are able to conceive the existence of truths that transcend our best capacities of proof, ascertainment, or verification. Moreover, he takes the instance of mathematics as a prime exhibit for anti-realism despite what would seem the inherent implausibility of any argument that confines mathematical truth to the compass of our best available proof-procedures or utmost computational powers. Here if anywhere there seems good reason to suppose (1) that the range of objective truths outruns our optimal capacity for proving, conceiving, or expressing them, and (2) that those truths decide the validity of our various well-formed (truth-apt) statements or theorems, rather than the other way around.[37] At least his approach has the virtue of posing these issues in their sharpest possible form and obliging his opponents to formulate their case with maximum care and precision so as to avoid falling into some well-laid anti-realist traps. Indeed it is the claim most often advanced on behalf of Dummett's pre-eminent status in current philosophical debate that he has managed to come up with a radical redefinition of the terms on which this longstanding dispute (i.e., between realism and anti-realism) must henceforth be conducted.

Of course one might interpret that claim as bearing only on certain rather technical or specialised issues in philosophy of language and logic, and hence as stopping well short of the extreme proposal that reality *just is* whatever we make of it according to the scope and limits of human perceptual, cognitive, or epistemic grasp. However this interpretation runs up against problems when it comes to Dummett's (so far as one can

tell) quite seriously meant talk about 'gaps in reality' and also those essays – like 'Bringing About the Past' – where he seems more than half-way convinced that changes in our present state of knowledge concerning past events can somehow influence (or even retroactively determine) the occurrence, non-occurrence, character, or outcome of those 'same' anterior events.[38] It is here that anti-realism in the Dummettian (analytic or logico-linguistic) mode comes closest to that strain of idealist thinking exemplified by the Oxford philosopher J. M. McTaggart whose influence Dummett readily admits in his own approach to these questions.[39] There is also a parallel with certain rather *outré* quantum-theoretical conjectures such as that of the astrophysicist John Wheeler who suggests – on the basis of laboratory-scale experiments to prove the existence of superluminal (faster-than-light) communication between pairs of remotely 'entangled' particles – that the same might apply to the retrocausal effect of momentarily switching a radio-telescope parameter and thus 'bringing about' some celestial event like a supernova at some billions of light-years' distance.[40] My point is that Dummett's 'technical' arguments in philosophy of language and logic have large (and quite drastically revisionist) implications for our thinking about issues in epistemology, ontology, and metaphysics. As regards their proper order of priority he maintains that this is the right way around and that logico-semantic considerations are our best guide to the settlement of issues in other, more contentious or less clearly demarcated regions of philosophical dispute.[41] All the same – as I have said – one may reasonably doubt whether Dummett's address to these matters is motivated solely (or chiefly) by his interest in sorting out the scope and limits of truth-talk in various contexts of enquiry or regions of discourse. Indeed one might go so far as to suggest that very often the metaphysical tail is wagging the logico-semantic dog, or that Dummett's more technical discussions of the realism/anti-realism issue are motivated in large part by his concern with questions such as that of the possible efficacy of prayer in deciding the as-yet unknown outcome of past events. (His example here involves the predicament of a father who prays that his son should not have been killed in a battle that has already taken place.)[42]

I am not making the claim that anti-realism in its current, Dummettian or logico-linguistic mode amounts to just a kind of technical camouflage for theological or metaphysical interests that dare not quite speak their name. After all it is a doctrine (or research-programme) that has not only captured the high ground of recent philosophical debate but succeeded in convincing a good many thinkers of an otherwise contrary (realist) persuasion that its arguments are sufficiently strong to require a very detailed and sophisticated effort of rebuttal.

Thus there is something inherently plausible about the basic anti-realist point, i.e., that if truth is conceived as objective (= recognition-transcendent) then *by very definition* it lies beyond our furthest powers of perceptual, cognitive, epistemic, or conceptual grasp. The standard test-case – at least for anti-realists – is that of mathematics where the argument goes that the realist is inevitably backing a loser since there seems no way that we could possibly have contact with (or epistemic access to) a realm of abstract entities such as numbers, sets, or classes which *ex hypothesi* transcend or exceed our capacity to comprehend them.[43] Hence the seeming paradox much exploited by sceptics and anti-realists: that we can *either* have mathematical truth realistically (objectively) conceived *or* mathematical knowledge within the limits of proof or computability but surely not both unless at the cost of embracing a Platonist conception whereby knowledge somehow links up with truth via some kind of sublimated (quasi-perceptual) means of access.

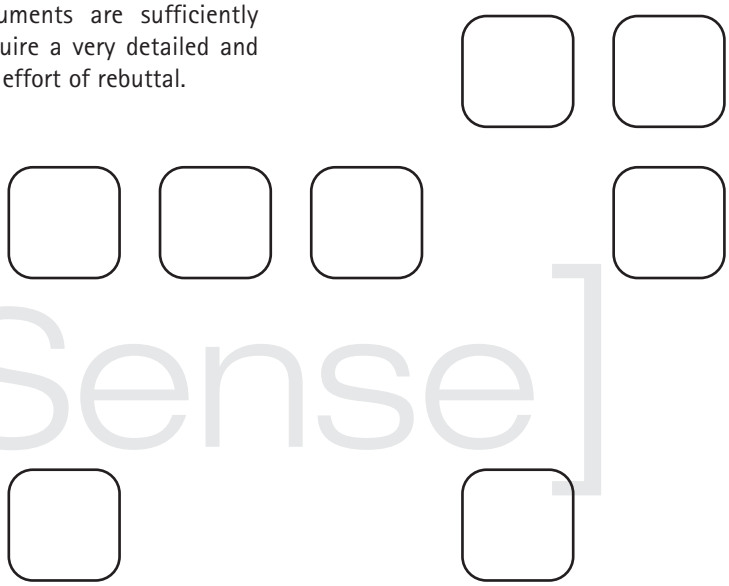

III

As I say, this line of argument is apt to strike one as possessing a knock-down philosophical force if taken on its own terms, i.e., on the assumption that these are the only alternatives and hence that realist (objectivist) truth in mathematics, logic, or the formal sciences cannot be conceived except as transcending – and *ipso facto* eluding – any knowledge we could possibly have of it. Yet it is likely to seem altogether less persuasive if one weighs it against the opposed considerations brought up by mathematical realists. Thus there is an irony about the fact that anti-realists have often claimed support from Gödel's incompleteness proof, that is, his demonstration that any system sufficiently complex to generate the axioms of elementary arithmetic will necessarily contain certain theorems which cannot themselves be proven within that system.[44] However – as exegetes like Penrose are quick to point out – this result, so far from counting against the existence of verification-transcendent truths, in fact lends weight to just the opposite (realist) conclusion, i.e., that we are capable of knowing that such truths exist despite their transcending the limits of formalised proof or computability.[45] Gödel himself put the case against a good many current anti-realist arguments when he wrote that 'mathematical intuition need not be conceived as a faculty giving an *immediate* knowledge of the objects concerned . . . Rather, they, too, may represent an aspect of objective reality, but, as opposed to sensations, their presence in us may be due to another kind of relationship between ourselves and reality'.[46] That is to say, the realist about mathematics need not be saddled with anything like the 'sublimated Platonist' conception of knowledge – the idea of our somehow having quasi-perceptual epistemic 'contact' with a realm of purely abstract entities – that is often foisted upon her by sceptics of various persuasion.[47] Moreover this alternative Gödelian view (taken up and developed by recent advocates of a rationalist-realist approach) manages to avoid some of the drastically counter-intuitive conclusions that result from Dummettian anti-realism when applied to particular cases.[48] Among them, for instance, is the absurdity of thinking that Fermat's Last Theorem – or the statement 'Fermat's Last Theorem is true' – was itself somehow neither true nor false until just that moment, after three centuries of failed efforts, when David Wiles traversed the last stage of his immensely complex and elaborate proof.[49]

No doubt it may be said that the proof was subject to challenge when first announced, then revised and strengthened in response to that challenge, and indeed might yet (quite conceivably) turn out to contain some further, as yet unnoticed weakness or logical flaw which casts doubt on its validity. However this objection is no more damaging to the mathematical realist's case than the similar argument brought against defenders of realism in the physical sciences. There it takes the form (as we have seen) of a sceptical meta-induction, or generalised 'argument from error', to the effect that most scientific theories to date have either been proved false or shown to hold good only within some restricted range of application, along with the various object-terms whose ontological standing was dependent on their role within those (nowadays discredited or superseded) theories. So the idea that we are now any better off in this respect – that our currently accredited theories are an exception to the general rule – must involve a high degree of epistemological hubris and also a failure, on the part of realist philosophers of science, to learn the most striking lesson offered by their historically-minded colleagues.[50] Yet it is precisely the realist's point – to repeat – that this argument itself *cannot but* have recourse to the conception of truth as transcending (and potentially falsifying) any particular thesis advanced at any stage in the history of scientific thinking to date. Thus it takes for granted the basic convergent-realist claim that theories and their associated object-terms may be subject to revision, qualification, or outright rejection on the strength of later (more adequate) evidential or theoretical-explanatory grounds.[51] After all it is no part of the realist's case to argue for our present state of scientific knowledge as secure against possible challenge or as having at last come out beyond any prospect of falsification. Indeed, as Nicholas Rescher points out, it is precisely this acceptance of the 'non-finality of science as we have it' – of the fact that even our most secure or well-established theories might always, in principle, be subject to challenge – that constitutes the realist's chief argument for the existence of objective, recognition-transcendent, or (at present) unverifiable truths.[52] Thus the standard sceptical meta-induction from past errors to the error-prone nature of all, including our current-best and future most

advanced states of scientific knowledge is an argument that the realist can turn back against the sceptic to powerful effect.

The case is rather different with mathematics since here any true proposition or valid proof must be taken to hold necessarily – or, in modal-logical parlance, across all 'possible worlds' – and therefore cannot be subject to disconfirmation as the result of some anomalous empirical result or some piece of conflicting evidence turned up in the subsequent course of enquiry. Jerrold Katz makes this point in a passage that also brings some useful clarification to the issue about Platonism and 'epistemic contact', so I shall take leave to quote it at length.

> The entire idea that our knowledge of abstract objects might be based on perceptual contact is misguided, since, even if we had contact with abstract objects, the information we could obtain from such contact wouldn't help us in trying to justify our beliefs about them. The epistemological function of perceptual contact is to provide information about which possibilities are actualities. Perceptual contact thus has a point in the case of empirical propositions. Because natural objects can be otherwise than they actually are (*non obstante* their essential properties), contact is necessary in order to discover how they actually are . . . . Not so with abstract objects. They could not be otherwise than they are . . . . Hence there is no question of which mathematical possibilities are actual possibilities. In virtue of being a perfect number, six must

> be a perfect number; in virtue of being the only even prime, two must be the only even prime. Since the epistemic role of contact is to provide us with the information needed to select among the different ways something might be, and since perceptual contact cannot provide information about how something must be, contact has no point in relation to abstract objects. It cannot ground beliefs about them.[53]

No doubt the anti-realist will protest that this simply begs the question with regard to the existence (or objective reality) of those various abstract items – numbers, sets, classes, etc., along with the range of true or false propositions concerning them – which he (the anti-realist) takes to 'exist' only in so far as they play some role in our present-best reasonings or proof-procedures. At which stage, perhaps, we should draw the conclusion that this a dispute beyond hope of settlement on any terms acceptable to both parties since it is one that involves such a sharp divergence of metaphysical views. Still the realist need not be stuck for an answer even if it is one that the anti-realist will routinely dismiss as buying into a naïve ('Platonist') metaphysics and a notion of our somehow having epistemic 'contact' with suchlike abstract entities, no matter how explicitly Katz, Gödel and others may have argued against that idea. Thus she can always point out that anti-realism leads to some downright bizarre claims, such as (to repeat) the idea that Fermat's Last Theorem possessed no objective truth-value until a proof was forthcoming. Or again, according to the anti-realist, it was neither true nor false that 311 successive iterations of the digit '1' constituted a prime number right up until the time when that fact emerged through the development of a computer programme with sufficiently powerful means of factorial analysis.

If these claims strike us as wholly implausible – as representing something like a *reductio ad absurdum* of the anti-realist case – then the same must apply to instances, like that of Goldbach's Conjecture, which involve well-formed and (on the face of it) truth-apt theorems but for which we lack computational means or any adequate proof-procedure that would decide their truth-value either way. For there is no reason – verificationist prejudice apart – to accord such instances special-case treatment and suppose that *just because* they remain unproven (and perhaps forever unprovable) *therefore* we are strictly enjoined to regard them as lacking such a value. Rather we should think – by analogy with those other kinds of case – that the issue concerning their truth and falsehood as a matter of objective (recognition-transcendent) mathematical fact is one that remains entirely unaffected by our present (or even our future-best) capacity to find it out. What Dummettian anti-realism amounts to, on this view, is an illicit extension of certain sceptical arguments as first applied to the methods and procedures of empirical enquiry – in particular Hume's problem about inductive reasoning – so as to encompass mathematics, logic, and the formal (axiomatic-deductive) sciences. Thus, for Dummett, it is crucially a matter of how we can justify talk of truth where such talk involves some delusive (objectivist) appeal to standards or criteria beyond those which we are enabled to grasp through our capacity to *recognise* the relevant truth-conditions and to *manifest* that knowledge in our various practices of formal reasoning. That is to say, just as Hume denied the validity of induction since we could never have demonstrative (logical) grounds for our belief in the existence of causal regularities in nature – such as were presupposed by any attempt to vindicate the claims of inductive warrant – so Dummett denies that we could ever have grounds for supposing mathematical or other kinds of truth to be verification-transcendent.

This comparison may appear less strained if one reflects on the striking resemblance between Hume's sceptical argument (i.e., that causal explanations always and inevitably go beyond the straightforward evidence of the senses) and Dummett's anti-realist proposal (i.e., that we venture onto perilous terrain if we suppose that truth-values can possibly transcend the limits of formal proof or empirical verification). What they both refuse to entertain, albeit on very different philosophical grounds, is the notion that we might have rational warrant for supposing certain statements to be true or false as a matter of the way things stand with respect to some given (whether abstract, physical, or real-world contingent) state of affairs, quite apart from any question concerning our sources of evidence or the scope and limits of our epistemic powers. Such is at any rate the basic realist position as defined by contrast with Dummett's type of logico-linguistic – though also, as I have said, metaphysically motivated – anti-realist argument. Indeed it is among the more curious features of this

whole debate that Dummett's way of framing the issue has so successfully managed to impose its preferential agenda and thereby steered discussion away from other, as one might think more central and substantive topics of concern.

Michael Devitt registers this sense of skewed priorities when he asks what rational justification there could possibly be for construing the issue about scientific realism in truth-theoretic terms and thence – through a further twist of anti-realist logic – as crucially involving our powers of linguistic or logico-semantic grasp. 'Realism', he writes,

> is an overarching empirical (scientific) theory or principle. It is initially plausible. It is supported by arguments that make no appeal to theories of language or understanding . . . . What firmer place could there be to stand than Realism, as we theorise in such undeveloped areas as those of language and understanding? In contrast, the poor state of theories in those areas, whether verificationist or not, makes them a bad place from which to start theorising, particularly in determining overarching principles about the nature of reality. To think otherwise is to put the cart before the horse.[54]

From this point of view the realist should reject Dummett's agenda, that is to say, his claim that the issue can best be treated as one concerning the existence (or non-existence) of recognition-transcendent truths, or of bivalent truth-values pertaining to statements of the 'disputed class'. To be sure, it is fundamental to the realist's case that Dummett's argument should *not* go through and that we *can* make sense of the contrary thesis, i.e., that our various well-formed and truth-apt (even if unverified or unverifiable) statements have their truth-value fixed – objectively so – by whether or not they correspond to the way things stand in reality. However she (the realist) will wish to go further and explain how we can none the less claim to have acquired knowledge of some such truths through various well-tried investigative methods and procedures. It is at this point – where metaphysical concerns yield ground to epistemological interests – that the argument is joined by other parties, among them advocates of the case for convergent realism (or inference to the best explanation) as the only means by which to make sense of advances in scientific knowledge to date.[55] Hence Devitt's thought that there is something strictly preposterous – a plain case of 'putting the cart before the horse' – about the notion that a theory (such as scientific realism) which enjoys such a vast range of corroborative evidence should be subject to doubt on the evidence of a relatively 'undeveloped' theory (such as Dummett's logico-linguistic approach) which exerts nothing like so strong a claim on our rational allegiance.

Of course any argument along these lines will fail to impress the convinced anti-realist for whom it is merely begging the question – i.e., the central issue as posed by Dummett – to take the (presumed) self-evidence of scientific progress as trumping the (presumed) highly fallible or dubious case from philosophy of language. No more will it persuade the van Fraassen-type constructive empiricist that his scruples are surely misplaced since all the evidence from scientific history to date points toward a different conclusion. That is to say, it lends weight to the convergent-realist claim that our received physical theories – e.g., with respect to atoms or subatomic particles – have typically advanced from a speculative stage, through a subsequent phase when such items acquired a crucial explanatory role yet when most physicists adopted an attitude of cautious (instrumentalist) reserve as concerned their objective reality, and thence to the point where those doubts became otiose with the advent of more refined observational or measurement techniques.

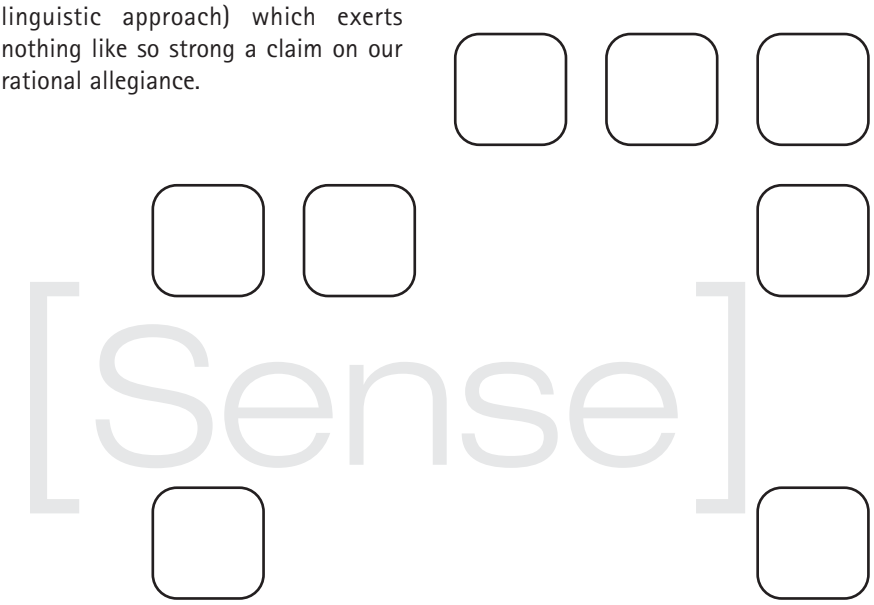

As I have said, van Fraassen would reject this account by arguing that such techniques – *just because* they are so refined or technologically advanced – can provide nothing like the probative warrant of direct 'naked eye' observation.[56] Still one may think it decidedly odd (another case of putting the cart before the horse) when van Fraassen draws his line for admission to the class of 'real' objects at the limit-point of plain, unaided human perceptual capacity. For this is to ignore a chief lesson from the history of science to date, namely that progress has most often come out through a break with the common-sense habit of relying on 'straightforward' perceptual self-evidence and a willingness to advance alternative theories and hypotheses.

These latter have ranged all the way from the most basic causal-explanatory conjectures – indispensable to science whatever Hume's (and van Fraassen's) sceptical thoughts on the matter – to the positing of certain as-yet unobservable objects (whether subatomic particles or planets) whose existence is deduced from their necessary role in resolving otherwise intractable problems and anomalies. Here again the typical pattern of development is from well-formed, truth-apt, but as-yet unverifiable (or unfalsifiable) hypotheses to theories so framed as to be capable of proof with some further – scientifically conceivable – advance in our means of testing them against the empirical evidence. However this is definitely *not* to maintain (like van Fraassen) that there is no going beyond the empirical evidence at any stage of scientific enquiry. For such a doctrine would preclude the very possibility of achieving any further advances of the kind that brought about the displacement of Ptolemaic by Galilean astronomy, or Newtonian by Einsteinian space-time physics, or pre-quantum by post-quantum conceptions of subatomic structure. That is, it would result in the arrest of scientific progress at whatever stage happened to mark this unfortunate relapse into naïve ideas of empirical self-evidence or the anthropomorphic (pre-scientific) notion that the limits of direct human perceptual acquaintance are the limits of attainable knowledge. Besides, as I have said, there is the further telling objection to van Fraassen's line of approach – one borne out by a vast range of neurophysiological and cognitive-psychological research – that what he takes as 'direct' sensory uptake is in fact no such thing but the product of various, immensely complex operations of perceptual processing.[57] Thus it is the merest of entrenched 'common-sense' prejudices that would attach more weight to the deliverance of (so-called) 'naked eye' perception than to the kinds of technologically enhanced observation made possible by sophisticated instruments whose workings (and whose possible defects, limits, or interference-effects) we are well placed to understand since, after all, they have been designed and constructed on established scientific principles. At least one may claim with good warrant that we now know more – with benefit of just those technologies – than was known when we had to rely on 'direct' sensory acquaintance or on comparative crude prosthetic devices like optical microscopes or telescopes.

In which case perhaps the undeniable subtlety, wit, and resourcefulness that van Fraassen deploys in support of his thesis should best be seen as something very like the impressive yet increasingly wire-drawn argumentation deployed by rearguard defenders of Ptolemaic astronomy against the new Copernican-Galilean cosmology. What has changed in the interim is that these problems have shifted from the first-order scientific terrain (where rival parties were divided with respect to the two 'world-systems' proposed by Ptolemy and Copernicus) to a meta-level dispute concerning the status of scientific knowledge in general and the existence – or otherwise – of truth-values that exceed the limits of empirical verifiability. This is the main point of convergence between van Fraassen's constructive-empiricist outlook and Dummett's anti-realist approach, despite their very different philosophical agendas, the one focused chiefly on epistemological issues and the other on issues in philosophy of language, logic, and metaphysics. Where they agree is in rejecting any talk of truth that exceeds the limits of empirical warrant (van Fraassen) or decidability according to our best available proof-procedures, sources of evidence, or means of verification (Dummett). Yet in both cases the argument runs up against a range of (to my mind) decisive objections. Among them is the fact that truth must play an indispensable role in any adequate characterisation of knowledge, and that one distinguishing mark of truth – except on the pragmatist conception of it as whatever is currently and contingently 'good in the way of belief' – is precisely its *not* being subject to the kinds of epistemic

limitation (or dependence on our current-best state of knowledge) entailed by such doctrines. Thus one is tempted to say that the *whole point* about truth, objectively conceived, is that it cannot be subject to the varying fortunes – including the chance of revision or downright disconfirmation – which always go along with epistemic conceptions like those of certainty, empirical warrant, 'truth' according to present best judgement, or even (at the limit) idealised rational acceptability. To suppose otherwise is simply to change the subject, or to find ways of redefining the truth-predicate so as to bring it safely back within the compass of humanly attainable knowledge.

IV

This strategy has exercised its strongest appeal among those most struck by the sceptical challenge in its latest (anti-realist) form, i.e., the idea that if truth is conceived in objectivist (recognition-transcendent) terms then *ex hypothesi* it cannot be known. And indeed there is no way around that sceptical argument if one accepts (1) that truth-values are epistemically constrained, (2) that warranted assertibility is the furthest we can get in such matters, and (3) that any thought of truth as transcending the limits of assertoric warrant is a thought that inevitably self-destructs on the manifest absurdity of claiming to know – to assert as a matter of truth – what exceeds our best means of proof or verification. It is not hard to see why anti-realism in this highly sophisticated logico-semantic guise has acquired such prominence in recent debate and spawned such a vast literature devoted to defending, strengthening, further refining, or (in some cases) trimming its claims so as to avoid any too direct conflict with realism as regards this or that specific area of discourse.[58] After all, it trades on the *prima facie* plausible idea that there *must* be something wrong – conceptually confused – about assertions of kind: 'I know statement x to be true [or false] even though I possess no means or method whereby to verify [or falsify] *x* and, what's more, no grasp of the conditions (i.e., those for warranted assertibility) under which I might come to recognise its truth-value and manifest my knowledge of them'. However the case looks far less plausible if one re-phrases the realist claim to read: 'I know that certain well-formed and truth-apt statements are *either* true *or* false – objectively so – despite my present and even (perhaps) despite anyone's future inability to verify or falsify those statements'. For it then becomes clear that the first way of putting the realist claim – embroiling it in patent absurdity or self-contradiction – simply begs the question since it takes for granted the anti-realist premise that truth is epistemically constrained and hence that the realist *cannot but* be making a strictly nonsensical statement.

However, as the second version makes clear, this is not at all what the realist has in mind since of course she rejects that premise outright (holding truth-values to be recognition- or verification-transcendent), and is therefore committed to nothing like the confusion so misleadingly foisted upon her by the anti-realist. At which point she can best turn the tables – though without any philosophic sleight-of-hand – and ask what further, more convincing justification the anti-realist can offer in support of a position which now looks to bear a much heavier burden of proof. Thus he will need to make good such claims as that Fermat's Last Theorem was neither true nor false until its proof was at last achieved, or that the truth-value of certain statements concerning remote astrophysical objects and events is determined by the scope and limits of human observation rather than decided – as the realist would have it – by astrophysical reality.

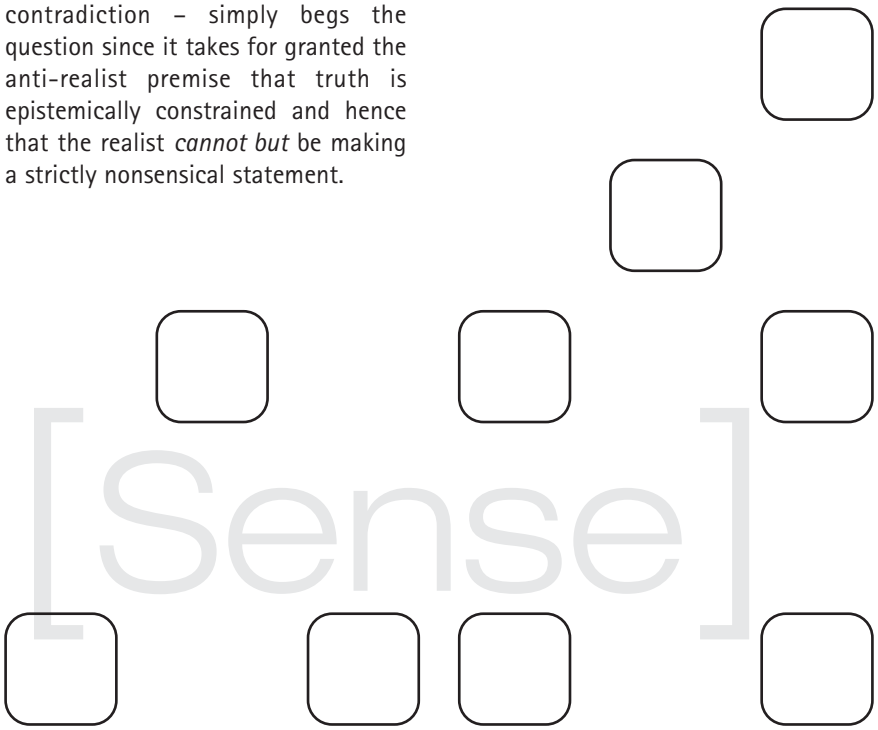

Even those of a marked anti-realist persuasion who have taken Dummett's lessons very much to heart quite often have trouble in going along with the consequences of his argument when spelled out in such explicit or case-specific terms. Thus some – Crispin Wright among them – have advanced various middle-way proposals which acknowledge the force of that argument with regard to any kind of full-strength 'metaphysical' realism while conserving a place for certain of our deep-laid realist intuitions as applied (say) to mathematics or the physical sciences.[59] However, as I have argued at length elsewhere, such efforts always end up *either* by endorsing the realist (objectivist) case in a form hedged about by various merely notional caveats and qualifying clauses *or* by falling back to a fairly standard version of the anti-realist line with just a few accommodating nods toward the kinds of realist objection noted above. The reason is plain enough: that there is simply no negotiating a midway or viable compromise solution with respect to those well-developed and conceptually precise areas of discourse – such as mathematics, logic, and the formal sciences – where any least concession to the view of truth as epistemically constrained or recognition-dependent is enough to constitute a repudiation of realism, albeit (very often) one that dare not quite speak its name.

So, for instance, when Wright puts forward his notions of 'superassertibility' and 'cognitive command' he is careful to specify the relevant criteria for statements of each type in terms that would appear to meet the realist's objection by building in additional constraints beyond those of (mere) assertoric warrant. 'Superassertibility' he defines as an attribute pertaining to any statement just on condition that 'some warrant for it would survive arbitrarily close scrutiny of its pedigree and arbitrarily extensive increments to or other forms of improvement of our information'.[60] For a discourse to exhibit 'cognitive command' is for statements of that discourse to meet the requirement that 'any difference of opinion will be such that there are considerations quite independent of the conflict which, if known about, would mandate withdrawal of one (or both) of the contending views'.[61] However these are still *epistemic* constraints, as can plainly be seen from such locutions as 'scrutiny of its pedigree', 'improvement of our information', and – lest 'quite independent of the conflict' be taken to lean too far in a realist direction – the crucial rider 'if known about'. Thus for all Wright's desire to accommodate the realist on the main points at issue with respect to certain such areas of discourse his approach still works out as an endorsement (albeit a somewhat queasy endorsement) of the anti-realist case. This emerges with particular clarity in his treatment of mathematics where Wright evinces a marked reluctance to go all the way with Dummett's constructivist, intuitionist, or proof-theoretic (as opposed to truth-based) conception yet conspicuously draws back from asserting any full-fledged realist commitment. Thus: 'in shifting to a broadly intuitionistic conception of, say, number theory, we do not immediately foreclose on the idea that the series of natural numbers constitutes a real object of mathematical investigation, which it is harmless and correct to think of the number theoretician as explaining'.[62] I can see no way of interpreting this oddly contorted sentence unless as a sop (more respectably: a source of reassurance) to the mathematical realist hedged around by various knowing asides – among them the adjective 'harmless' – designed to placate those of Dummettian persuasion who will no doubt bridle at any such concessions to the adversary camp.

There is a similar unresolved tension in recent attempts by other philosophers of mathematics to come up with some middle-ground formulation that would save realist appearances while yielding no hostages to objectivist (and hence, on their own terms, sceptical) fortune. These involve the idea of a 'humanised Platonism' which, unlike its 'sublimated Platonist' counterpart, brings the whole issue intelligibly down to earth in those various mathematical practices, reasonings, and warranted proof-procedures that constitute truth so far as it can possibly be known.[63] On this account truth is 'conceptually structured' – and hence within epistemic reach – yet still somehow capable of offering guidance (or correcting our erroneous judgements) when we are disposed to get things wrong. What prevents us from seeing this is an unfortunate attachment to the kind of sublimated Platonist conception which equates truth with something that stands intrinsically above and beyond our best powers of epistemic grasp. Hence the colourful analogy drawn by Alex Miller in his debunking estimate of what gives rise to the objectivist delusion (along with the equally disabling sceptical backlash) in philosophy of mathematics. 'In our pre-theoretical

thinking', he writes,

> we have a perfectly healthy desire for a degree of independence between our judgements and the facts which those judgements are capable of tracking. When we do philosophy, this healthy desire becomes sublimated into an *unhealthy philosophical conception* of what this independence has to consist in. So just as Gustav Mahler's perfectly healthy respect for women becomes sublimated into an unhealthy syndrome known as the Virgin Mary complex, our own perfectly healthy desire for a measure of independence between the knower and what is known becomes sublimated into the idea that the properties which the judgements of the knower cognitively access have to be conceptually unstructured.[64]

We can best get over this unhealthy fixation – so the argument goes – if we cease the vain hankering for objective truths that could somehow (impossibly) be accessed quite apart from our means of coming to know them. Rather we should see that mathematical knowledge is in no way compromised or rendered less secure by its dependence on our various reasonings, reckonings, or established proof-procedures. That is to say – and here Miller takes his cue from John McDowell – the whole misbegotten congeries of problems around truth, knowledge, and scepticism begins with that delusive ('sublimated') Platonist conception of truth which assigns it to a realm of absolute ideal objectivity beyond any epistemic contribution on the knower's part.[65] Where the hard-line realist goes wrong is in supposing that 'we can think of our judgements about the instantiation of a property as capable in principle of tracking or cognitively accessing the facts about its instantiation only if the property in question is conceptually unstructured'. On the humanised Platonist account, conversely, we can 'think of ourselves as tracking or cognitively accessing the facts about the instantiation of conceptually *structured* properties'.[66]

Miller has his own differences with McDowell as regards the precise working-out of this approach. Still he shares McDowell's basic conviction that the only way around the 'problem of knowledge' with regard to mathematics and other truth-apt areas of discourse is one that makes room for the conceptual structuring of everything that falls within their remit and which thus restores truth to the compass of humanly attainable knowledge. However this solution just won't work, as becomes clear from McDowell's often tortuous attempts to explain how one can have a fully adequate measure of objectivity (i.e., an account of how truth might always come apart from best judgement or even from the standard of idealised rational warrant) along with an epistemic approach that restricts truth-values to the range of statements for which we possess some demonstrable means of proof or verification.[67]

Hence McDowell's (in my view) somewhat desperate proposal that we should go back to Kant for a viable alternative to the way these issues have been treated in the wake of logical empiricism, i.e., an approach that makes room for the joint and strictly inseparable contributions of Kantian 'receptivity' and 'spontaneity'.[68] Thus we are to think that these latter are really just *faute de mieux* terms of art which denote on the one hand the mind's responsiveness to objective (non-mind-dependent) inputs or sources of knowledge and on the other its inbuilt 'spontaneous' power to cognise or apprehend such truths. All the same, McDowell cautions, they should properly be thought of as aspects or components of one and the same knowledge-constitutive capacity. Where the error comes in is with the dualist notion (also much encouraged by Kant) that the business of philosophy is somehow to explain how two such heterogeneous 'faculties' as sensuous intuition and conceptual understanding can be brought together through a faculty of judgement whose ultimate source is the power of productive imagination, itself defined as as 'a blind but indispensable function of the soul, without which we should have no knowledge whatsoever, but of which we are scarcely ever conscious'.[69] McDowell sees clearly that the travails of much analytic philosophy from the logical positivists and logical empiricists down have resulted from this bad Kantian inheritance, one that fixes an insuperable gulf between truth (or reality) and our knowledge of it and which then goes various intricate and ultimately self-defeating ways around in solving the problem thus produced. Much better start out from Kant's alternative ideas of 'receptivity' and 'spontaneity' since these make room for a non-dualist conception whereby we can at last 'dismount from the seesaw' since the two terms can be taken as referring to the self-same cognitive or epistemic capacity which brings truth back within the compass of humanly attainable knowledge.

However McDowell's argument breaks down on the fact that he is still very firmly seated on the Kantian seesaw, and one whose oscillations cannot be damped by switching from talk of 'intuitions' and 'concepts' to talk of 'receptivity' and 'spontaneity'. That is, such talk still leaves it a mystery (one much exploited by sceptics and anti-realists) how we could ever gain knowledge of truths that none the less obtained quite apart from our evidential sources, i.e., our best methods of formal proof or empirical verification. 'If we restrict ourselves to the standpoint of experience itself', McDowell writes,

> what we find in Kant is precisely the picture I have been recommending: a picture in which reality is not located outside a boundary that encloses the conceptual sphere . . . . The fact that experience involves receptivity ensures the required constraint from outside thinking and judging. But since the deliverances of receptivity already draw on capacities that belong to spontaneity, we can coherently suppose that the constraint is rational; that is how the picture avoids the pitfall of the Given.[70]

Yet this can scarcely be supposed to resolve the problem – one that McDowell inherits as much from Kant as from the doctrines of logical positivism or logical empiricism – if one considers the extreme contortions of phrasing (and the wrenchings of logical thought) forced upon him by the effort to reconcile the claims of objective, mind-independent truth and attainable knowledge. Thus it is hard to make sense of his idea that thinking and judgement are somehow 'constrained' by that which lies 'outside' their spontaneous grasp – through a power of receptivity that is subject to constant checks and corrections from the external world – while that constraining influence is nevertheless thought of as 'draw[ing] on capacities that belong to spontaneity'. Confusion is worse confounded – or so it seems to me – when McDowell talks about 'reality' as that which is 'not located outside a boundary that encloses the conceptual sphere'. For in that case reality *just is* whatever falls within the scope and limits of our perceptual, cognitive, or epistemic grasp and cannot be conceived as potentially transcending our knowledge of it.

U

What we are getting here, in effect, is a warmed-over ('analytic') version of the history of German idealism after Kant. Such was the debate between, on the one hand, 'subjective idealists' like Fichte who purported to follow Kant's doctrine to its ultimate conclusion by treating reality as a construct or projection of our egological concepts and categories and, on the other, 'objective idealists' like Schelling who sought to maintain some 'external' (mind-independent) check on those same concepts and categories.[71] What we are also getting is a vague adumbration of some quasi-Hegelian synthesis that would emerge on the far side of all those vexing Kantian antinomies and occupy a standpoint above and beyond their inherently limiting or partial perspectives. However this standpoint turns out to be no such thing but to take us straight back onto the ground of subjective idealism, albeit hedged about by various quasi-objectivist caveats and scruples. Thus, according to McDowell,

> [i]t can be difficult to accept that the Myth of the Given is a myth . . . . It can seem that we are retaining a role for spontaneity but refusing to acknowledge any role for receptivity, and that is intolerable. If our activity in empirical thought and judgement is to be recognisable as bearing on reality at all, there must be external constraint. There must be a role for receptivity as well as spontaneity, for sensibility as well as understanding. Realising this, we come under pressure to recoil back into appealing to the Given, only to see over again that it cannot help. There is a danger of falling into an interminable oscillation.[72]

Still it is far from clear that McDowell has managed to dismount from the seesaw whose oscillations Kant set going through his heroic though ultimately failed attempt to reconcile the twin doctrines of 'empirical realism' and 'transcendental idealism'. Indeed one could write the history of much post-1950 (that is to say, post-logical-empiricist) work in the broadly analytic tradition as a series of projects aimed toward mending the Kantian rift between phenomenal intuitions and concepts of understanding but always – inevitably – running up against the same root dilemma.[73] What has united these movements despite and across some otherwise large differences of view is their shared premise that objectivist (alethic) realism must surely give rise to scepticism by placing truth by very definition beyond our utmost cognitive grasp. Whence the whole range of alternative proposals – from Dummett's anti-realist agenda to response-dispositional theories and Wright's sundry variations on the theme – that seek to bring truth back within the sphere of human cognitive or intellectual grasp. Yet their upshot is chiefly to exacerbate the problem (and induce yet further swings of the Kantian seesaw) by adopting an epistemic approach which, no matter how nuanced or conceptually refined, fails to uphold the crucial distinction between truth or veridical knowledge on the one hand and, on the other, such fallback notions as 'cognitive command', 'superassertibility', 'best judgement', or 'idealised rational warrant'.

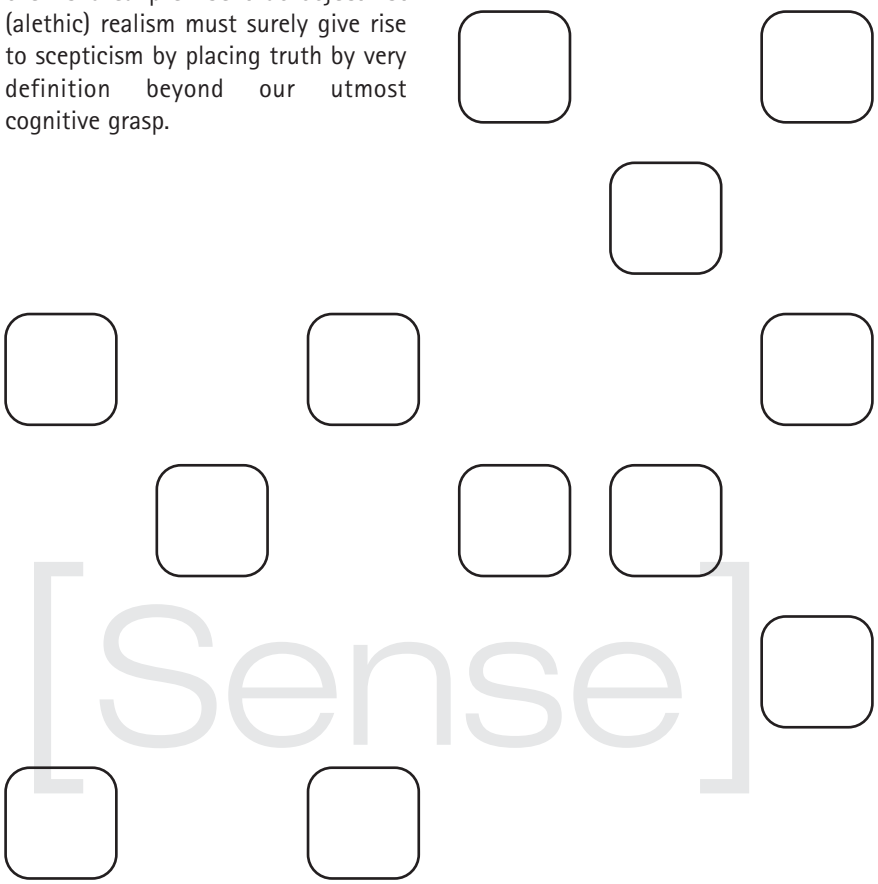

It seems to me that this problem must remain strictly insoluble so long as philosophers persist in confusing metaphysical with epistemological issues, i.e., questions concerning the structure and content of truth with questions concerning our various kinds and degrees of epistemic justification. No doubt this will again be thought to beg the question against anti-realism since it is just Dummett's point that the two sorts of issue are inextricably bound up together. As we have seen, what leads him to adopt that approach is a range of logico-semantic considerations with their chief source in Frege and their upshot in a metaphysical doctrine with far-reaching epistemological consequences. Thus, according to Dummett, by far the best hope of achieving greater clarity about this issue is to come at it *via* debates in philosophy of language and logic where we are on much firmer conceptual ground than when forwarding large (and inherently contentious) claims about the progress of the physical sciences to date or realism as a matter of inference to the best, most rational explanation. However we should here recall Devitt's argument to contrary effect, i.e., that anti-realism puts the epistemo-logico-linguistic horse before the scientific cart by taking its cue from a relatively 'underdeveloped' area of discourse (philosophical semantics) and attaching a wholly disproportionate weight to the kinds of problem that result.[74] At any rate there is something distinctly awry about a theory that purports to resolve these issues – even to prevent them from getting off the ground – while in fact blocking their solution at every turn. What anti-realism chiefly serves to show, as I have argued, is the impossibility of carrying its premises through to a credible conclusion and the fact that we can make rational sense of advances in the physical and formal sciences only on a realist or alethic (truth-based) approach to the various issues involved.

References

1 See especially Michael Dummett, *Truth and Other Enigmas* (London: Duckworth, 1978), *The Logical Basis of Metaphysics* (Duckworth, 1991), and *The Seas of Language* (Oxford: Clarendon Press, 1993); also Michael Luntley, *Language, Logic and Experience: the case for anti-realism* (Duckworth, 1988); Neil Tennant, *Anti-Realism and Logic* (Oxford: Clarendon Press, 1987) and *The Taming of the True* (Oxford: Oxford University Press, 1997).

2 Dummett, *Truth and Other Enigmas* (op. cit.), p. xl.

3 Dummett, *The Seas of Language* (op. cit.), p. 468.

4 Dummett, *Elements of Intuitionism*, 2nd edn. (Oxford: Clarendon Press, 2000).

5 See Dummett, *Truth and Other Enigmas* (op. cit.); also *Frege and Other Philosophers* (Oxford: Clarendon Press, 1991).

6 Dummett, *Truth and Other Enigmas* (op. cit.), p. 155.

7 Dummett, *The Seas of Language* (op. cit.), p. 75.

8 I take this example from Scott Soames, *Understanding Truth* (Oxford: Oxford University Press, 1999).

9 For some pertinent discussion, see C.J. Misak, *Verificationism: its history and prospects* (London: Routledge, 1995).

10 Bas van Fraassen, *The Scientific Image* (Oxford: Clarendon Press, 1980); also *Laws and Symmetry* (Clarendon, 1989).

11 van Fraassen, *The Scientific Image* (op. cit.), p. 8.

12 Ibid, p. 12.

13 van Fraassen, 'Empiricism in the Philosophy of Language', in Paul Churchland and Clifford Hooker (eds.), *Images of Science: essays on realism and empiricism, with a reply from Bas C. van Fraassen* (Chicago: University of Chicago Press, 1985; p. 255.

14 See for instance J.L. Aronson, 'Testing for Convergent Realism', *British Journal for the Philosophy of Science,* Vol. 40 (1989), pp. 255-60; Aronson, R. Harré and E. Way, *Realism Rescued: how scientific progress is possible* (London: Duckworth, 1994); Richard Boyd, 'The Current Status of Scientific Realism', in Jarrett Leplin (ed.), *Scientific Realism* (Berkeley & Los Angeles: University of California Press, 1984), pp. 41-82; Gilbert Harman, 'Inference to the Best Explanation', *Philosophical Review*, Vol. 74 (1965), pp. 88-95; Peter Lipton, *Inference to the Best Explanation* (London: Routledge, 1993).

15 M. Gardner, 'Realism and Instrumentalism in Nineteenth-Century Atomism', P*hilosophy of Science*, Vol. 46 (1979), pp. 1-34; Ian Hacking, *Representing and Intervening: introductory topics in philosophy of science* (Cambridge: Cambridge University Press, 1983); Mary Jo Nye, *Molecular Reality* (London: MacDonald, 1972); J. Perrin, *Atoms*, trans. D.L. Hammick (New York: van Nostrand, 1923).

16 See especially Paul Churchland, 'The Ontological Status of Observables: in praise of the superempirical virtues', in Churchland and Hooker (eds.), *Images of Science* (op. cit.); also Christopher Norris, 'Anti-Realism and Constructive Empiricism: is there a (real) difference?' and 'Ontology According to van Fraassen: some problems with constructive empiricism', in *Against Relativism: philosophy of science, deconstruction and critical theory* (Oxford: Blackwell, 1997), pp. 167-95 and 196-217.

17 See for instance Rodolfo Llinas and Patricia Churchland, *The Mind-Brain Continuum: sensory processes* (Cambridge, MA: M.I.T. Press, 1996); A.D. Milner, *The Visual Brain in Action* (Oxford U.P., 1995); David Rose and Vernon G. Dobson (eds.), *Models of the Visual Cortex* (Chichester: Wiley, 1985); J.Z. Young, *Philosophy and the Brain* (Oxford U.P., 1987).

18 See Hilary Putnam, 'Language and Reality', in *Mind, Language, and Reality* (Cambridge: Cambridge University Press, 1975), pp. 272-90; p. 290; also Boyd, 'The Current Status of Scientific Realism' (op. cit.).

19 See entries under Note 14, above.

20 For further discussion of these and other such cases, see Stathis Psillos, *Scientific Realism: how science tracks truth* (London: Routledge, 1999).

21 See for instance Dummett, *On Immigration and Refugees* (London: Routledge, 2001).

22 Dummett, *The Logical Basis of Metaphysics* (op. cit.), p. 7.

23 Ibid, p. 7.

24 See for instance Bertrand Russell, *Our Knowledge of the External World as a Field for Scientific Method in Philosophy* (London: Allen & Unwin, 1914).

25 For a useful conspectus, see Richard Rorty (ed.), *The Linguistic Turn: recent essays in philosophical method* (Chicago: University of Chicago Press, 1967).

26 The issue about truth-value links receives some informative and shrewd discussion in Bernhard Weiss, *Michael Dummett* (Chesham: Acumen, 2002).

27 Bernard Williams, 'What Was Wrong with Minos?', in *Truth and Truthfulness: an essay in genealogy* (Princeton, N.J.: Princeton University Press, 2002), pp. 149-71.

28 Thucydides, *History of the Peloponnesian War*, trans. Rex Warner (Harmondsworth: Penguin, 1954).

29 Williams, *Truth and Truthfulness* (op. cit.), p. 163).

30 See Note 22, above.

31 Williams, *Truth and Truthfulness* (op. cit.), p. 168.

32 For a range of views, see Joyce Appleby, Lynn Hunt, and Margaret Jacob, *Telling the Truth About History* (New York: Norton, 1994); Richard Campbell, *Truth and Historicity* (Oxford: Oxford University Press, 1992); Richard Evans, *In Defence of History* (London: Granta Books, 1997); Christopher Norris, *Truth and the Ethics of Criticism* (Manchester: Manchester University Press, 1994); Paul Ricoeur, *History and Truth*, trans. Charles A. Kelbley (Evanston, IL: Northwestern University Press, 1965).

33 Williams, *Truth and Truthfulness* (op. cit.), p. 167.

34 Ibid, p. 167.

35 For further arguments to similar effect, see Michael Devitt, *Realism and Truth*, 2nd edn. (Oxford: Blackwell, 1986); Gerald Vision, *Modern Anti-Realism and Manufactured Truth* (London: Routledge, 1988);

36 Dummett, 'Can an Effect Precede its Cause?', 'Bringing About the Past', and 'The Reality of the Past', in *Truth and Other Enigmas* (op. cit.), pp. 319-32, 333-50 and 358-74.

37 See especially Jerrold J. Katz, *Realistic Rationalism* (Cambridge, MA: M.I.T. Press, 1998).

38 See Note 36, above.

39 McTaggart, John, *Philosophical Studies*, ed. S.V. Keeling (London: Longmans, 1934).

40 J.A. Wheeler, 'Delayed Choice Experiments and the Bohr-Einstein Dialogue'. Paper presented at the joint meeting of the American Philosophical Society and the Royal Society, London, June 5th, 1980. See also F. Selleri, 'Wave-Particle Duality: recent proposals for the detection of empty waves', in W. Schommers (ed.), *Quantum Theory and Pictures of Reality: Foundations, interpretations, and new aspects* (Berlin: Springer Verlag, 1989), pp. 279-32; J.A. Wheeler and W.H. Zurek (eds.), *Quantum Theory and Measurement* (Princeton, N.J.: Princeton University Press, 1983)

41 See especially Dummett, *The Logical Basis of Metaphysics* (op. cit.).

42 See Dummett, 'Bringing About the Past' (Note 36, above).

43 See Paul Benacerraf, 'What Numbers Could Not Be', in Benacerraf and Hilary Putnam (eds.), *The Philosophy of Mathematics: selected essays*, 2nd edn. (Cambridge: Cambridge University Press, 1983), pp. 272-94; also Michael Detlefson (ed.), *Proof and Knowledge in Mathematics* (London: Routledge, 1992); W.D. Hart (ed.), *The Philosophy of Mathematics* (Oxford: Oxford University Press, 1996); Hilary Putnam, *Mathematics, Matter and Method* (Cambridge University Press, 1975).

44 See Kurt Gödel, 'On Formally Undecidable Propositions of *Principia Mathematica* and Related Systems', trans. B. Meltzer (New York: Basic Books, 1962); also Ernest Nagel and James Newtman, *Gödel's Theorem* (London: Routledge & Kegan Paul, 1971) and S.G. Shanker (ed.), *Gödel's Theorem in Focus* (London: Routledge, 1987).

45 Roger Penrose, *Shadows of the Mind:a search for the missing science of consciousness* (London: Vintage Books, 1994).

46 Kurt Gödel, 'What is Cantor's Continuum Problem?', in Benacerraf and Putnam (eds.), *The Philosophy of Mathematics* (op. cit.), pp. 470-85; p. 484.

47 See Benacerraf, 'What Numbers Could Not Be' (op. cit.); also John Divers and Alexander Miller, 'Arithmetical Platonism: reliability and judgement-dependence', *Philosophical Studies*, Vol. 95 (1999), pp. 277-310 and Miller, 'Rule-Following, Response-Dependence, and McDowell's Debate with Anti-Realism', *European Review of Philosophy*, Vol. 3 (1998), pp. 175-97.

48 See Notes 44 and 45, above; also Katz, *Realistic Rationalism* (op. cit.).

49 For an informative 'popular' account, see Simon Singh, *Fermat's Last Theorem: the story of a riddle that confounded the world for 358 years* (London: Fourth Estate).

50 See especially Larry Laudan, 'A Confutation of Convergent Realism', *Philosophy of Science*, Vol. 48 (1981), pp. 19-49.

51 See entries under Note 14, above.

52 Nicholas Rescher, *Scientific Realism: a critical reappraisal* (Dordrecht: D. Reidel, 1987), p. 61.

53 Katz, *Realistic Rationalism* (op. cit.), pp. 36-7.

54 Michael Devitt, *Realism and Truth* (op. cit.), p. 284.

55 See entries under Note 14, above.

56 van Fraassen, *The Scientific Image* (op. cit.).

57 See entries under Note 17, above.

58 See Note 1, above; also – for a critical review of these developments with extensive bibliography – Christopher Norris, *Truth Matters: realism, anti-realism, and response-dependence* (Edinburgh: Edinburgh University Press, 2002).

59 Crispin Wright, *Truth and Objectivity* (Cambridge, MA: Harvard University Press, 1992).

60 Ibid, p. 48.

61 Ibid, p. 103.

62 Ibid, p. 5.

63 See Miller, 'Rule-Following, Response-Dependence, and McDowell's Debate with Anti-Realism' (op. cit.).

64 Ibid, p. 178.

65 See especially John McDowell, 'Intentionality and Interiority in Wittgenstein', in K. Puhl (ed.), *Meaning Scepticism* (Berlin: de Gruyter, 1991), pp. 148-69 and 'Meaning and Intentionality in Wittgenstein's Later Philosophy', *Midwest Studies in Philosophy*, Vol. 17 (1992), pp. 40-52.

66 Miller, 'Rule-Following, Response-Dependence, and McDowell's Debate with Anti-Realism' (op. cit.), p. 178.

67 See especially McDowell, 'Wittgenstein on Following a Rule', *Synthèse*, Vol. 58 (1984), pp. 325-63.

68 McDowell, *Mind and World* (Cambridge, MA: Harvard University Press, 1994). For further discussion see Norris, 'McDowell on Kant: redrawing the bounds of sense' and 'The Limits of Naturalism: further thoughts on McDowell's *Mind and World*', in *Minding the Gap: epistemology and philosophy of science in the two traditions* (Amherst, MA: University of Massachusetts Press, 2000), pp. 172-96 and 197-230.

69 Immanuel Kant, *Critique of Pure Reason*, trans. N. Kemp Smith (London: Macmillan, 1964).

70 McDowell, *Mind and World* (op. cit.), p. 41.

71 For a well-informed survey of these developments, see Frederick C. Beiser, *The Fate of Reason: German philosophy from Kant to Fichte* (Cambridge, MA: Harvard University Press, 1987).

72 McDowell, *Mind and World* (op. cit.), pp. 8-9.

73 See Norris, *Minding the Gap* (op. cit.); also Michael Friedman, *A Parting of the Ways: Carnap, Cassirer, and Heidegger* (La Salle, IL: Open Court, 2000).

74 Devitt, *Realism and Truth* (op. cit.).

Notes On Contributors

Will Cartwright

is lecturer in philosophy at the University of Essex. He has an academic background in both law and philosophy, and specialises in a range of areas including legal and moral philosophy, medical ethics and theory of responsibility. His publications include 'Killing and Letting Die' in the *British Medical Bulletin* (52:2, 1996), and 'Equal Access to Health Care: a Problematic Ideal' in Tom Sorrell (ed), *Health Care, Ethics* (London, Routledge, 1998).

David Papineau

is professor of philosophy at King's College London. He completed undergraduate degrees in both mathematics and philosophy, before completing his doctorate at Cambridge. He has taught in a range of top philosophy departments in the UK and elsewhere, specialising in philosophy of science, philosophy of mind and psychology, and epistemology. His books include *Reality and Representation* (Oxford: Blackwell, 1987), *Philosophical Naturalism* (Oxford: Blackwell, 1993) and *Thinking About Consciousness* (Oxford: OUP, 2002). He was also one of the speakers at the first RJP annual conference in the spring of 2002.

John Collins

is lecturer in philosophy at UEA (Norwich), and has previously taught at Birkbeck College, London, University of Warwick, and APU (Cambridge). His current interests centre on the relations between generative linguistics and the philosophies of mind and language. He also has a longstanding interest in the concept of truth and the philosophies of logic and mathematics. John has published numerous papers in journals such as *Mind, Synthese, Philosophy and Phenomenological Research, Erkenntnis, etc.* He is currently writing a book on the concept of logical form.

Paul Sperring

is joint editor of the RJP.

Christopher Norris

is Distinguished Research Professor in Philosophy at the University of Cardiff. He has also taught at the University of Duisburg and the University of Wales Institute of Science and Technology. His wide range of interests includes philosophy of language and mind, metaphysics and metaethics, and issues in epistemology and science. His long list of publications includes *Resources of Realism: Prospects for 'Post-Analytic' Philosophy* (London: Macmillan, 1997) and *Against Relativism: Philosophy of Science, Deconstruction and Critical Theory* (Oxford: Basil Blackwell, 1997). He has published previously in the RJP.

Kathy Behrendt

is a lecturer in Philosophy at New College, University of Oxford. She is originally from Canada, and studied both there and in Oxford. Her main research interests are in the areas of personal identity, self-consciousness and memory, and she has published a number of articles in these areas.

Notes For Contributors

Content

We welcome articles on any area in philosophy. Papers may be broad or narrow in their focus (for instance a discussion of the mind/body problem, or an analysis of Hume's treatment of causation in the *Enquiry*). We would particularly encourage contributions which reflect original research on the following philosophical themes: epistemology, metaphysics, philosophy of religion, ethics, philosophy of mind, philosophy of science, political philosophy, religious ethics; and texts, such as: *The Republic, The Nicomachean Ethics, The Meditations, An Enquiry Concerning Human Understanding, Beyond Good and Evil, On Liberty, Existentialism and Humanism, The Problems of Philosophy, Language Truth and Logic*.

The articles should be around 3000-4000 words.

Style

The language used in the articles should be as non-technical as possible whilst preserving the richness of the arguments. Where technical terms are unavoidable they should be explained and examples offered.

Notes should be presented as endnotes. Textual references should be presented in the following format: Barry Stroud, *Hume* (London: Routledge, 1977), 77-91.

Presentation

Articles should be written in *Word* (any version).

Contributions

Articles for this journal are currently written by a panel of philosophers from a variety of universities in Britain, Australia and the United States, whose work is edited by the journal's editorial board. We invite unsolicited contributions from philosophers working in any field. The contributions should be submitted via email attachment to rjp@rutc.ac.uk

Copyright

Contributors

Richmond upon Thames College

is a large further education college located in
Twickenham, West London
offering 16-19 students one of the widest
curriculum choices in the country.

Last year, we came top of all
London further education colleges
in the Times league tables
and we are proud of our reputation
for achieving excellent results year after year.

We are well known nationally for
our high quality staff, excellent student support systems
and the inclusive education we offer to all our students.

If you would like to find out more about us -
please contact our Course Information Unit on

020 8607 8305 / 8314

or visit our website on

www.rutc.ac.uk

Richmond upon Thames College

[Subscribing]

to the *RJP*

The RJP comes out in Autumn, Spring and Summer. To subscribe you need to select the appropriate price from the table and complete the mailing information at the bottom of the form. All prices include post and packaging. The bottom half of the form should then be detached and sent with a cheque made payable to Richmond upon Thames College to the address below:

RJP Subscriptions
Philosophy Department
Richmond upon Thames College
Egerton Road
Twickenham
London TW2 7SJ
United Kingdom

Please allow one week for delivery in the UK, and two weeks for the rest of the world.

Annual Subscription : Current Rates

Institutional Subscriber in the UK	£33.00
Individual Subscriber in the UK	£18.00
Institutional Subscriber from the rest of the EU	63.00 Euros
Individual Subscriber from the rest of the EU	40.50 Euros
Institutional Subscriber outside the EU	$67.50 US
Individual Subscriber outside the EU	$45.00 US

[subscription]

Your name: ..... Your organisation's name (if appropriate) .....

Address .....

..... Post code/Zip code.....

Telephone number (including full international dialling code) .....

Email address .....

I enclose a cheque for the amount of .....to purchase an annual subscription

Signed ..... Date .....

[Philosophy]

The Richmond Journal of **Philosophy**

Richmond Journal of Philosophy
Philosophy Department
Richmond upon Thames College
Egerton Road
Twickenham
Middlesex
TW2 7SJ
United Kingdom

Tel: 020 8607 8000
Fax: 020 8744 9738
Email: rjp@rutc.ac.uk
www.rutc.ac.uk/rjp

Richmond upon Thames College

ISSN: 1477-6480

Design & Production Marketing for Education 01282 612222 [ref 26097 - 11/03]