

VATS, SETS, AND TITS*

1. At the beginning of Iris Murdoch's novel *The Black Prince* the narrator mentions four episodes from his story that might serve as suitable starting points and comments, 'There are indeed many places where I could start.'¹ I feel a bit like that.

I could start with my title. Some explanation is certainly called for! I could start with a parenthetical remark that Hilary Putnam makes in *Reason, Truth and History*. Referring to his celebrated argument that we could not be brains in vats, Putnam says, 'This argument first occurred to me when I was thinking about a theorem in modern logic, the "Skolem-Löwenheim Theorem".'² I could start with the very idea of critical self-conscious reflection on one's most basic beliefs about oneself and one's environment. I could start with my ultimate concern: transcendental idealism. Any of these would give my essay a suitable steer—except, perhaps, the first, which would just allay curiosity. But the curiosity is liable to be distracting, so I shall in fact start there.

For as long as I can remember, the abbreviation that I have used in my note-taking for 'things in themselves' has been 'tits'. And, since this essay will be concerned with comparisons and contrasts between three paradigms, in one of which things in themselves are the counterparts of vats and sets in the other two, I found this title irresistible. But I note that the English word 'tits' has a number of meanings, even when not being used in what the dictionary would classify as a vulgar way. It can denote a kind of songbird of course; it can be used as a pejorative slang word to denote horses or young women; as a Scots word, it can denote twitches or tugs; and it can denote a gentle kind of knock, as in the phrase 'tit for tat'. In due course there will, I hope, be something agreeably apposite about this miscellany of very phenomenal associations.

* This essay draws together ideas that I have expounded elsewhere, principally in Moore (1996) and Moore (1997), Ch. 7, §1.

¹ Murdoch (1975), p. 21.

² Putnam (1981), p. 7.

2. To return to *Reason, Truth and History*. In Chapter 1 of that book Putnam considers the following scenario: a human brain is kept alive in a vat of nutrients and is manipulated by scientists in such a way as to give the subject the hallucinatory experience of living a perfectly normal life with a perfectly normal body. This scenario presents an obvious sceptical challenge: how do I know that I am not in this predicament? Putnam rises to the challenge by arguing that, at least in a sufficiently drastic version of the scenario,³ its protagonist—let us call him Brain—cannot so much as entertain the thought that he is in this predicament, whatever other thoughts he might be able to entertain; which means that anyone who does entertain the thought that he is in this predicament thereby believes that very thought; which means that, when *I* entertain the thought that *I* am in this predicament, I believe my thought; which provides me with a way of meeting the sceptical challenge.

This argument has generated a large and fascinating literature.⁴ One common objection to the argument is that, even if it gives me an assurance that I am *not* in the relevant predicament—an assurance, in other words, that my thought that I am just a brain in a vat (in sufficiently drastic circumstances) cannot be true—it does so only at the price of showing that I may not fully grasp this thought. *I may not know what this thought comes to at the level of things in themselves*. For, if Putnam's argument is correct, then Brain might entertain a thought *of the same type* as mine, with an assurance *of the same type* that his thought cannot be true, little realizing that, really, what his thought comes to is some very complex conjecture about a possible configuration of things in a vat which, really, he is just a brain in. And the new sceptical challenge is:

³ The reason for this qualification is that a good deal depends on what sort of causal contact there is, and has been in the past, between the brain and its environment: see Smith (1984). Henceforth in this section I shall tend to take this qualification for granted, though in §3 below it will be crucial and I shall once again make it explicit.

⁴ See e.g. Brueckner (1986); Sacks (1989), Ch. 3; Wright (1994), to which Putnam replies in Putnam (1994), §8; and Forbes (1994).

how do I know that I am not in a predicament analogous to *this*? As Crispin Wright puts it,

[the] real spectre to be exorcized concerns the idea of a thought *standing behind* my thought that I am not a brain in a vat, in just the way that my thought that it *is* a mere brain in a vat would stand behind the thought... of an actual brain in a vat that ‘I am not a brain in a vat’... What I should really like would be an assurance... not just that most of what I think is actually true... but that I am on to the right categories in terms of which to depict the most general features of the world and my place in it.⁵

Anyone sympathetic to Putnam’s argument is liable to say that this new sceptical challenge can be met in precisely the same way as the original one. Brain can no more entertain thoughts about what ‘stands behind’ his thoughts, or about ‘the right categories’ in terms of which to depict ‘the most general features of the world’, than he can entertain thoughts about brains and vats: that is, real brains and real vats.⁶

But those who see a genuinely new challenge here have two possible responses. The first of these is to distinguish between concepts such as that of a general feature of the world or that of a thing in itself, on the one hand, and concepts such as that of a brain or that of a vat, on the other, and to insist that Brain *can* entertain thoughts involving concepts of the former kind. This might be because concepts of the former kind are *a priori*, whereas concepts of the latter kind are empirical, so that possessing concepts of the former kind does not require

⁵ Wright (1994), pp. 239 – 241, adapted from the first-person plural to the first-person singular, his emphasis.

⁶ Putnam himself responds to the new sceptical challenge in this way: see Putnam (1994), pp. 286 – 288.

the same sort of interaction with one's environment as possessing concepts of the latter kind does.⁷

The second response is to grant Brain, and ultimately each of us, an insight into the possibility that there is some radical difference between how things are in themselves and how they appear which does not take the form of a normal thought at all; indeed, which might even be beyond our powers of expression.⁸

How different is this second response from the first? That depends on what counts as a 'normal' thought and on whether the insight in question *is* in fact supposed to be beyond our powers of expression. If a 'normal' thought is simply a thought that we can express, which is arguably just another way of saying that a 'normal' thought is simply a thought, and if, in accordance with that, the insight *is* supposed to be beyond our powers of expression, then the two responses are very different. If a 'normal' thought is a thought that involves concepts that are unavailable to Brain, because possessing them involves a certain sort of interaction with one's environment, and if the insight is *not* supposed to be beyond our powers of expression, then there may not be any difference at all. However that may be, both responses are gestures in the direction of a kind of transcendental idealism, whereby our 'normal' thoughts answer merely to how things appear, not to how they are in themselves. In Kant, from whom of course such transcendental idealism derives, the contrast between thoughts that are 'normal' in this way and

⁷ Cf. Kant (1998), A85/B117. Note that, in order for this response to succeed, only a few concepts need to be of the former kind. And if they include the concept of *analogy*, then the response will not only succeed but succeed in a way that is very powerful, in the sense that it will significantly enlarge the range of sceptical thoughts that are available: consider, for instance, my thought that some being is to me in this or that respect as I am to Brain. See further Moore (1996), §5, and the discussion there of Nagel (1986). Cf. also my very use of the word 'analogous' in stating the new sceptical challenge above.—For a very interesting discussion relevant to the question whether the concept of the *self* is of the former kind, and therefore possessable by Brain, see Madden (unpublished), where Madden raises some concerns about whether Brain can indeed think about, or refer to, himself. For remarks relevant to the possibility that he cannot, see Evans (1982), pp. 251 – 252. Also very relevant are Kant's remarks on self-consciousness at Kant (1998), B157.

⁸ Cf. Moore (1997), Ch. 7, §1.

thoughts that are not is the contrast between thoughts that have content and thoughts that do not, where this in turn is the contrast between thoughts that involve intuitions and thoughts that do not. Kant holds that thoughts of the former kind answer to how things are empirically: how they appear to beings with the relevant intuitions. These, for Kant, are the only thoughts that can constitute knowledge. But he allows for thoughts of the latter kind as well, abstract unverifiable conjectures about how things are in themselves.⁹ And whatever we make of this contrast and its attendant metaphysics, it looks as though *that* is the sort of position that we shall end up in if we try to resist Putnam's argument to the last.¹⁰

3. To give these considerations somewhat sharper focus I shall now recast them in terms of three paradigms, which I shall call the Vat Paradigm, the Set Paradigm, and the Tit Paradigm.

The Vat Paradigm: I remarked above that Putnam's argument requires a 'sufficiently drastic' version of his scenario, but I did not elaborate.¹¹ Putnam himself envisages a universe in which *all* sentient creatures are brains in a vat being tended by automatic machinery that is programmed to give them a

⁹ See e.g. Kant (1998), Bxxv – xxvii, and Kant (1996), Pt I, Bk II, Ch. II, §VIII. John McDowell, in the opening section of McDowell (1994), denies that Kant allows for thoughts of the latter kind: 'abnormal' thoughts, in the terms that I have been using. Commenting on the famous passage in which Kant declares that thoughts without content are empty (Kant (1998), A51/B75), McDowell writes, 'For a thought to be empty... would be for it not really to be a thought at all, and that is surely Kant's point; he is not, absurdly, drawing our attention to a special kind of thoughts, the empty ones,' (pp. 3 – 4). But that is precisely what Kant is doing, or at least what he takes himself to be doing: see e.g. Kant (1998), A253 – 254/B309.

¹⁰ To be sure, 'it looks as though' is the operative phrase. There are all sorts of twists and turns in the dialectic which I have ignored and which would have to be considered in any full discussion of these matters. Some philosophers think that Putnam's argument can be resisted at a much earlier stage: see e.g. Nagel (1986), Ch. V, §2. (Note that on p. 73 of Nagel's discussion, he considers a version of the second response to Putnam's extended argument.) For further reservations about how much Putnam's argument achieves see Lewis (1984), pp. 233 – 236.

¹¹ See above, footnote 3.

collective hallucination.¹² As it happens, less drastic versions of his scenario would have suited his purposes just as well. But there are versions that are less drastic still, for which his argument certainly fails. The Vat Paradigm is one of these. It concerns a human brain that has only *recently* been envatted and is being manipulated by scientists in such a way as to make the subject think that nothing untoward has happened. The subject in this case—let us call him Cerebrum—can certainly entertain the thought that he is in that predicament. If he does entertain this thought, and somehow reassures himself that it is false, then he is wrong.

The Set Paradigm: This concerns a set theorist—let us call him Georg—who uses standard set-theoretical terminology,¹³ but under a non-standard interpretation. Georg's interpretation is elementarily equivalent to the standard interpretation,¹⁴ in other words it makes precisely the same sentences come out true; but it differs from the standard interpretation in being a countable sub-interpretation of it. (It is the Skolem-Löwenheim theorem, the theorem to which Putnam refers in his parenthetical remark, that guarantees the existence of such an interpretation.) Georg is oblivious to the possibility of an interpretation extending his in the way in which the standard interpretation does. But, because his interpretation is elementarily equivalent to the standard interpretation, the sentences that he holds true, at least in so far as he is good at what he does, are the sentences that *are*

¹² Putnam (1981), p. 6.

¹³ The formal language in which he works, and in terms of which this terminology is defined, is the first-order language whose sole extralogical constant is \in .

¹⁴ For current purposes I am simply taking for granted that there is such a thing as 'the standard interpretation'. Such an assumption is by no means philosophically innocuous, of course. But to address it would be a quite separate undertaking.

true, under the standard interpretation. Among these is the sentence, ‘There are uncountably many sets of finite ordinals.’ There is a construction which, under Georg’s interpretation, constitutes a proof of this sentence, even though his interpretation has only countably many sets in its domain altogether.¹⁵

The Tit Paradigm: This concerns a subject—let us call him Noumenon—who views the world through native spectacles of some metaphorical kind which he can never take off. He has no knowledge of how things are in themselves, only of how they appear through the spectacles. But worse: he cannot even entertain or express thoughts about how things are in themselves. What he can do, however, is to achieve an insight, beyond his powers of expression, into the possibility that he is subject to precisely such limitations.¹⁶

These three paradigms have an important feature in common. In all three cases there is a subject whose thinking, in a certain respect, is sensitive only to a limited aspect of how things are in that respect: what I shall call the subject’s ‘phenomenal bubble’.¹⁷ But there are also some important differences between the three paradigms. Cerebrum, unlike either Georg or Noumenon, is a victim of systemic error. Though his thinking is *sensitive* only to his phenomenal bubble, it is *answerable* to more than that. When he thinks, ‘I am at the post office,’ he is thinking something that would be true only if he were at the post office, that is

¹⁵ For further details and discussion see Moore (2001), Ch. 11.

¹⁶ This is obviously meant to call to mind—I shall put it no more strongly than that—the metaphysical picture that Kant paints in Kant (1998). (One reason for not putting it any more strongly than that is the point I made at the end of §2: Kant allows for thoughts about things in themselves.)

¹⁷ For example, Georg’s thinking, in respect of what sets are like, is sensitive only to what sets in the domain of his non-standard interpretation are like. Henceforth I shall normally take the qualification about the relevant ‘respect’ for granted

only if things beyond his phenomenal bubble were a certain way. By contrast, Georg's thinking and Noumenon's thinking are each sensitive and answerable to the same thing: each is both sensitive and answerable to the thinker's phenomenal bubble. Neither Georg nor Noumenon need be involved in any error at all. In Georg's case, if not perhaps in Noumenon's, another subject can have thoughts *of the same type* as his that are answerable to how things are beyond his (Georg's) phenomenal bubble. For instance, I can think that there are uncountably many sets of finite ordinals. But although my set-theoretical thought is of the same type as one of Georg's, it does not have the same *content*.¹⁸ That is why my thought and Georg's thought, despite being answerable to aspects of set-theoretical reality that differ in a way that might have been expected to confer different truth-values on them, are both in fact true. It is somewhat like the situation in which I watch a Wimbledon ball boy fumble a ball that has been gently lobbed to him, while Roger Federer watches his next opponent hit an unreturnable cross-court volley, and each of us thinks, quite correctly, 'I might have done that.' For that matter, it is somewhat like the situation in which I am at my desk, and Cerebrum is hallucinating being at his desk, and each of us thinks, quite correctly, 'I am *not* at the post office.' The difference is that, in the Set Paradigm, the relevant element of perspective is not located in anything as circumstantial as the use of the first person: it pervades the entire discourse. (And, of course, Georg, provided that he has arrived at his thought on sound mathematical grounds, is not *accidentally* correct in his thinking, in the way in which Cerebrum is.) Finally, the crucial difference between the Set Paradigm and the Tit Paradigm is that Noumenon, unlike Georg, can achieve an insight, albeit beyond his powers of expression, into the possibility that there is just such an element of perspective pervading everything he says and thinks; that everything he

¹⁸ For amplification of this distinction between a thought's type and its content see Moore (1997), pp. 9 – 11.

says and thinks is answerable only to his phenomenal bubble, not to things in themselves.

Corresponding to these three paradigms are three epistemological claims that we might make, which I shall refer to as vat-scepticism, set-scepticism, and tit-scepticism. Vat-scepticism is the claim that, for all I know, I am in Cerebrum's situation; set-scepticism is the claim that, for all I know, I am in Georg's situation; and tit-scepticism is the claim that, for all I know, I am in Noumenon's situation.

Now it is natural to assimilate ordinary philosophical scepticism concerning the truth of one's most basic beliefs about oneself and one's environment to vat-scepticism. It is natural, in other words, to think that the target of such scepticism is the possibility that one is a victim of systemic error of the sort that afflicts Cerebrum. And Putnam's argument, as we have noted, does nothing to rebut vat-scepticism. What it does do, however, if successful, is to call into question the assimilation, by showing how limited vat-scepticism is. In philosophical terms, Cerebrum's situation is not particularly outlandish. One would not have to have an especially philosophical cast of mind to think that vat-scepticism could *not* be rebutted; indeed, that it was true. But if Putnam is right, the more drastic possibilities of concern to philosophical sceptics are closer, in various critical respects, to the Set Paradigm than they are to the Vat Paradigm: in particular, Brain's predicament is closer, in these respects, to Georg's predicament than to Cerebrum's. (Hence Putnam's remark about the Skolem-Löwenheim Theorem.)

To put it in these terms, however, just seems to reinforce the original objection to Putnam's argument, which was that it answers one sceptical challenge only by presenting a new one. The new challenge can now be seen as a variation on *set*-scepticism. For suppose that Putnam is right and I am not a victim of systemic error of any drastic kind in my thinking about the external world. How do I know that this is not just because my thinking is answerable to nothing but my phenomenal bubble,

in a way of which I have not the least idea? Call this variant of set-scepticism set-like-scepticism.

Now set-scepticism itself—never mind, for the time being, any such non-mathematical variant of it—can certainly be rebutted. This is something that I have argued elsewhere, as indeed has Putnam, to whom I am indebted.¹⁹ The argument, in summary, runs as follows.

The Argument Against Set-Scepticism: Set-scepticism is the claim that, for all I know, I am in Georg's situation, which entails that, for all I know, my thinking about sets is never thinking about *all* sets, but only ever about those in the domain of my own restricted interpretation of the language of set theory. But even in entertaining the thought that I am in that situation, I believe it. For it is itself a thought about *all* sets.

The set-sceptic will of course caution that these references of mine to 'all' sets may themselves have limited scope. But what does this mean? That they may not really be references to *all* sets? Yet that is precisely what they are!

Is there an analogue of this argument serving to rebut set-like-scepticism? Apparently so. It runs as follows.

The Argument Against Set-Like-Scepticism: Set-like-scepticism entails that, for all I know, my thinking about the external world is never answerable to the external world in *all* its aspects, but only ever to my phenomenal bubble. But even in entertaining the thought that I am in that situation, I believe it. For it is itself a thought that is answerable to the external world in *all* its aspects.

¹⁹ Moore (2001), Ch. 11, esp. §3; and Putnam (1983).

Does this argument really serve to rebut set-like-scepticism? It may very well do, but arguably only at the price of showing that set-like-scepticism is not the scepticism that really concerns us; or in other words, that the assimilation of Brain's predicament to Georg's needs to be questioned. For Brain's predicament, however much like Georg's it may be, and in particular however much more like Georg's it may be in certain critical respects than it is like Cerebrum's, is nevertheless different from Georg's in other respects that are just as critical.

To amplify. The argument against set-like-scepticism is very similar to, if not identical with, the extension of Putnam's argument that we considered in the previous section. But when we considered that extension, we also considered two responses to it. And precisely what these responses both did, in effect, was to call into the question the assimilation of Brain's predicament to Georg's. The first response suggested that they are unlike each other in the following respect: not *all* of Brain's thoughts are answerable only to his phenomenal bubble. The second response, at least in its more radical forms, conceded that Brain's predicament is like Georg's in that respect—Brain's thoughts are indeed all answerable only to his phenomenal bubble—but fastened on a different difference. It suggested another shift of paradigm in fact: from the Set Paradigm to the Tit Paradigm. For the basic idea was that, even if I cannot *think* that my thoughts are all answerable only to my phenomenal bubble, still I can achieve an insight of sorts, albeit beyond my powers of expression, into the possibility that they are.²⁰ And if I *am* able to

²⁰ It is worth noting parenthetically that these are by no means the only reasons for resisting the assimilation of Brain's predicament to Georg's. Someone might argue, for instance, that there is nothing more to sets than what we think about them, in contrast to the things that Brain thinks about. Again, relatedly, someone might argue that there is nothing more to how a set theorist interprets the language of set theory than which sentences from that language he or she holds true, which means that the Set Paradigm is incoherent, or in other words that Georg's predicament is not a genuine possibility. But not even Putnam would say this about Brain's predicament: see Putnam (1981), pp. 7 – 8.

achieve such an insight, then we seem left with a final, unanswered sceptical challenge: a variation, this time, on *tit*-scepticism.

Or is ‘sceptical challenge’ the right description? ‘Tit-scepticism’ itself is just a label. It is a substantive question whether the claim to which that label attaches (namely, that, for all I know, I am in Noumenon’s situation) deserves to count as a species of *scepticism* at all. Might it not be better viewed as a gesture towards a familiar and redoubtable philosophical doctrine: transcendental idealism? Which of course brings us back to the position that we were in at the end of the previous section.

4. I am in no doubt that we have just witnessed one of the main impulses towards a radical version of transcendental idealism. The idea that I have some kind of insight into the impossibility of my thinking about things in themselves, and into the corresponding limitations of that which I *can* think about, is a very natural destination for the train of thought that departs from Putnam’s own original argument.²¹

Is it an attractive destination? My own view is that it is not. If we are to board this train at all—and I think there are good reasons for doing so—then I believe we should alight at an earlier stage, earlier even than the stage of conceding that, *for all I know*, such transcendental idealism is true. Come to that, I believe we should alight before acceding to a less radical version of transcendental idealism whereby the insight that I have into my own limitations is one that I can express, and the limitations are limitations only to that part of my thinking which is ‘normal’ in some

²¹ Putnam himself is not averse to the idea that there is a route from his argument to a kind of transcendental idealism: see e.g. Putnam (1981), pp. 60 – 64. Cf. also Sacks (1989), Ch. 3 (though Sacks construes transcendental idealism in a more epistemological way than I have been doing—as something very like *tit*-scepticism in fact).

robust sense.²² I believe we should alight at the argument against set-like-scepticism.²³

I cannot argue for this now. To do so would require saying considerably more than I have been able to say in this essay about why we should go even that far. But I shall close by saying a little about why I do not think that we should go all the way to the more radical version of transcendental idealism, even though there is a strong temptation to do so.

Part of the reason why there is a strong temptation to do so, it seems to me, is that we do have an inexpressible insight, activated by Putnam's argument, into what it is for that to which our thinking is answerable to be how we think it is; and, if we try to express this insight, then we are liable to circumscribe that to which our thinking is answerable and treat it as some kind of phenomenal bubble. We are liable to endorse the more radical version of transcendental idealism. But while this explains our temptation, it does not in any way vindicate it. The fact is, we cannot treat that to which our thinking is answerable as some kind of phenomenal bubble, for a reason famously articulated by Wittgenstein: 'in order to be able to draw a limit to thought, we should have to find both sides of the limit thinkable (i.e. we should have to be able to think what cannot be thought).'²⁴ (This is reminiscent of the argument against set-like-scepticism, the point at which I have already said I think we should be stationing ourselves: I cannot entertain the thought that my thinking is answerable only to my phenomenal bubble without having a thought that is answerable to more than my phenomenal bubble.)

But still, the transcendental idealist may say, if we really do *have* the inexpressible insight, and thereby share with Noumenon the crucial feature that distinguishes him from Georg, what is to stop us from simply alluding to that fact and thus assimilating our own case to the Tit

²² See again the material at the end of §2 above: Kant's own transcendental idealism is of this less radical kind.

²³ This means, in particular, that we should resist the attempts above to dissociate Brain's predicament from Georg's.

²⁴ Wittgenstein (1961), p. 3.

Paradigm? And if we do that, shall we not be *en route* to the more radical version of transcendental idealism after all?

Well, unless we actually try to express the insight, not just allude to the fact that we have it, we shall *not* have assimilated our own case to the Tit Paradigm. To share with Noumenon the crucial feature that distinguishes him from Georg, it is not enough that we have an inexpressible insight; obviously not. It is not even enough that we have an inexpressible insight that we are tempted to treat as an insight into the possibility that our thinking is answerable only to our phenomenal bubble. We need to have an insight that *really is* an insight into that possibility. And—quite apart from the fact that any insight that really was an insight into that possibility would *eo ipso* be expressible—if there is no such possibility, then there can be no such insight. So the most we can do is to acknowledge our own temptation to endorse the more radical version of transcendental idealism. And that clearly falls short of actually endorsing it. It even falls short of conceding its coherence.

If I am right that we should not endorse any such transcendental idealism, and if I am right about how far short of that we should allow this train of thought to take us, what does this mean, finally, about things in themselves? It certainly removes one substantial reservation that we might have about whether we can so much as think about things in themselves. But we should not conclude without further ado, as we might be tempted to, that things in themselves are just ordinary middle-sized dry goods and their like. It is possible, for instance, that the phrase ‘things in themselves’ is best viewed as syncategorematic, so that ‘thinking about things in themselves’ means something like ‘thinking about things in a way that is totally free of perspective, whether cultural, biological, or of any other kind’, in which case we may have to conclude that things in themselves ‘are’—to the extent that it is appropriate even to talk in these terms—the fundamental particles of physics or something of that sort.²⁵

²⁵ Cf. Moore (1997), Ch. 4, §4.

Still, even if things in themselves should not be thought of as including ordinary middle-sized dry goods, or the various episodes that involve them, we have some kind of reassurance that they can legitimately be regarded as being *of a piece with* ordinary middle-sized dry goods and the various episodes that involve them—such as songbirds, horses, and gentle knocks, things of the sort that can be denoted by the English word ‘tits’.

REFERENCES

- Brueckner, Anthony (1986), ‘Brains in a Vat’, in *The Journal of Philosophy* **83**
- Evans, Gareth (1982), *The Varieties of Reference*, ed. John McDowell (Oxford: Oxford University Press)
- Forbes, Graeme (1995), ‘Realism and Skepticism: Brains in a Vat Revisited’, in *The Journal of Philosophy* **92**
- Kant, Immanuel (1996), *Critique of Practical Reason*, trans. Mary J. Gregor, in Immanuel Kant, *Practical Philosophy*, trans. and ed. Mary J. Gregor (Cambridge: Cambridge University Press)
- Kant, Immanuel (1998), *Critique of Pure Reason*, trans. and ed. Paul Guyer and Allen W. Wood (Cambridge: Cambridge University Press)
- Lewis, David (1984), ‘Putnam’s Paradox’, in *Australasian Journal of Philosophy* **62**
- Madden, Rory (unpublished), ‘Could a Brain in a Vat Self Refer?’
- McDowell, John (1994), *Mind and World* (Cambridge, Mass.: Harvard University Press)
- Moore, A.W. (1996), ‘Solipsism and Subjectivity’, in *European Journal of Philosophy* **4**
- Moore, A.W. (1997), *Points of View* (Oxford: Oxford University Press)
- Moore, A.W. (2001), *The Infinite*, 2nd edn (London: Routledge)
- Murdoch, Iris (1975), *The Black Prince* (Harmondsworth: Penguin)

- Nagel, Thomas (1986), *The View From Nowhere* (Oxford: Oxford University Press)
- Putnam, Hilary (1981), *Reason, Truth and History* (Cambridge: Cambridge University Press)
- Putnam, Hilary (1983), 'Models and Reality', in Paul Benacerraf and Hilary Putnam (eds), *Philosophy of Mathematics: Selected Readings* (Cambridge: Cambridge University Press)
- Putnam, Hilary (1994), 'Comments and Replies', in Peter Clark and Bob Hale (eds), *Reading Putnam* (Oxford: Basil Blackwell)
- Sacks, Mark (1989), *The World We Found: The Limits of Ontological Talk* (London: Duckworth)
- Smith, Peter (1984), 'Could We Be Brains in a Vat?', in *Canadian Journal of Philosophy* **14**
- Wittgenstein, Ludwig (1961), *Tractatus Logico-Philosophicus*, trans. D.F. Pears and B.F. McGuinness (London: Routledge & Kegan Paul)
- Wright, Crispin (1994), 'On Putnam's Proof that We Are Not Brains in a Vat', in Peter Clark and Bob Hale (eds), *Reading Putnam* (Oxford: Basil Blackwell)