

Hilary Putnam (1926–)

JOHN HEIL

The Philosophical Lexicon contains the following entry for “hilary”:

hilary, n. (*from* hilary term) A very brief but significant period in the intellectual career of a distinguished philosopher. “Oh, that’s what I thought three or four hilaries ago.” (Dennett 1987: 11)

The entry makes reference to Hilary Putnam’s penchant for changing his views, even completely reversing himself on central themes.

What are we to make of this inconstancy? Emerson, in a famous but widely misquoted passage wrote:

A foolish consistency is the hobgoblin of little minds, adored by little statesmen and philosophers and divines. With consistency a great soul has simply nothing to do. (Emerson 1940: 152)

Philosophers, especially technically adroit philosophers, often adopt a position then stick with it no matter what. If you are clever, you can find ways around almost any objection. Defending a cherished thesis can be like defending an old friend against a charge of dishonesty: your own honor as well as your friend’s is at stake. In philosophy, however, candor trumps constancy. It is to Putnam’s credit that he has been willing to allow his views to evolve as they will even when this leads him in surprising directions.

Commendable as it is, this kind of intellectual forthrightness puts pressure on anyone setting out to summarize Putnam’s views. Not only has Putnam written on a wide variety of issues, but his take on those issues has shifted, often dramatically. There is not one Putnam, but many. In what follows I have selected from among the available Putnams those that seem to me to have had the most immediate philosophical impact. Inevitably, I have had to leave out much that is interesting. I shall not discuss any of Putnam’s important technical work in the philosophy of mathematics, logic, and the philosophy of science. (His views on these and many other topics can be found in Putnam 1975a, 1975b, and 1983.) I shall focus – selectively – on three domains in which Putnam has had considerable influence over philosophy as it is now practiced: philosophy of language, philosophy of mind, and metaphysics.

Philosophy of language

Suppose you utter the sentence “That’s a banana” meaning to indicate a banana. What is it about you that makes it the case that your utterance, accompanied perhaps by a gesture, indicates a banana? One natural response is that the pertinent feature of you that make it the case that your utterance concerns a banana is your own state of mind. When you deliver the utterance, your mind is focused in a certain way (on bananas!), and this mental focusing is what gives your linguistic utterance its significance. I *understand* your utterance as indicating a banana when the utterance triggers in me a comparable state of mind. This state of mind constitutes my *grasp* of your utterance’s significance.

Proponents of the traditional *ideational* model of meaning appeal to ideas or mental images. As you utter the sentence, you entertain a banana image. The image secures a connection between your words and the world. Imagery is not essential, however. A meaning might be a kind of definition you carry around inside your head, a recipe that tells you when to apply a particular term. (Let us ignore a regress problem that seems to undercut any such view: if we understand terms only by possessing a definition, what enables us to understand the terms constituting that definition?) Think of images and mental recipes as mechanisms for fixing the extension of terms. (The extension of a term is the set of objects it designates. The English word “water,” for instance, designates a kind of stuff. When you use this word, you designate that stuff.)

Approaching the topic from this direction invites us to distinguish sharply between *meaning* and *reference*. The *meaning* of “water” includes only elements that you the speaker can grasp. Competent speakers need know nothing of the chemical composition of water; they may be ignorant of the nature of the stuff designated by the term “water.” Even those with an intimate knowledge of the constitution of water are rarely in a position to apply this knowledge to determine when a liquid substance is or is not water. If we eliminate specialized knowledge as a requirement for knowing the meaning of “water,” we are left with the idea that the term means a clear, colorless, tasteless liquid found in oceans, ponds, and rain puddles. Your grasp of this meaning is what enables you to use the term “water” correctly. The meaning, understood as a kind of recipe or rule grasped by speakers, is what connects the word “water” to the stuff, water.

Wittgenstein inaugurated a sustained attack on views of this kind in the period between the two world wars (see Wittgenstein 1953). According to Wittgenstein, meaning is determined by social and contextual factors. Your utterance means what it does, not because of a state of mind that lies behind the utterance, but because you produce the utterance as a member of a particular community of language users. In this community words are seamlessly integrated with actions and with interactions with non-linguistic states of affairs. To understand the meaning of a word is to understand how it figures in the practices – linguistic and otherwise – of members of your community (see WITTGENSTEIN). This understanding is ultimately grounded in your capacity to *engage* in the pertinent practices.

Putnam’s approach to meaning can be seen as involving an articulation and extension of Wittgenstein’s insight. Putnam begins with an attack on the familiar distinction between meaning and reference. We learn to use the term “water” at an early age.

Later we learn that the stuff to which “water” refers is H_2O . In this way we discover empirically what water *is*. This discovery, however, sheds light as well on what our term “water” *means*. In speaking of water we mean to be speaking about stuff like *this* stuff (here we point to some water). We may know nothing of the stuff’s hidden nature, but we use “water” to designate the stuff with the nature of *this* stuff, whatever that might be. The stuff in question is, as we now know, H_2O . So water’s *being* H_2O is part of the *meaning* of “water.”

Twin Earth

To reinforce this point, Putnam invites us to imagine a distant planet that precisely resembles Earth in almost every respect (see Putnam 1975b: ch. 12). Its continents are arranged exactly as our continents are arranged, its inhabitants speak languages precisely resembling languages spoken on Earth. Were you instantaneously transported to this planet, you would notice no differences at all. Inhabitants of the planet who live in a place they call “America” refer to their planet as “Earth,” but to avoid confusion let us dub it Twin Earth. Twin Earth differs from Earth in just one respect: the colorless, tasteless, odorless liquid stuff that fills Twin Earth oceans, rivers, ice trays, and fish tanks is not H_2O , but a different substance, XYZ. Although XYZ superficially resembles H_2O , it possesses a very different chemical constitution. In other respects, however, Twin Earth precisely resembles Earth down to the last detail.

Suppose that you utter the sentence, “I’ll have a glass of water, please.” Your utterance concerns water and, assuming it occurs in appropriate circumstances (you are not rehearsing for a play, for instance, or making a philosophical point), you are issuing a request for a glass of water. Imagine, now, that your twin on Twin Earth produces an exactly resembling utterance. Your twin’s utterance does not concern water, nor does your twin request a glass of water. Water is H_2O , and the stuff called “water” on Twin Earth is not H_2O , but XYZ. We might say that your twin’s utterance of “water” concerns *twin* water; your twin is requesting a glass of *twin* water.

You and your Twin Earth counterpart may be as alike as you please (leaving aside the fact that the chemical constitution of your respective bodies will be importantly different!); the images running through your mind could precisely resemble the images running through your twin’s mind; your feelings could be the same. Yet your utterance and your twin’s appear to have different meanings. Given the intrinsic similarities between you and your twin, the meanings of words you utter must be determined by something other than your intrinsic makeup. You and your twin may be entirely ignorant of the chemical constitution of what you each call “water.” Indeed, English speakers (and their counterparts on Twin Earth) used the term for generations before anyone was in a position to appreciate that water was a particular sort of chemical compound. Earlier, we regarded this as a good reason to suppose that the meaning of words must be limited to what speakers can individually grasp. But this conception of meaning as strongly distinguished from reference is precisely what Putnam’s Twin Earth thought experiment challenges. If we think of the meanings of our terms as what fixes the extension of those terms – where the extension of a term is just the stuff or set of objects designated by the term – then we must give up the idea that meanings are like pictures or recipes we carry around inside our heads and consult when we apply words to objects.

The division of linguistic labor

English speakers use the terms “beech” and “elm” to designate species of tree. If you are like me, however, you would be hard pressed to say how beeches and elms differ and utterly unable to distinguish a beech from an elm in the wild. Does this mean that, for English speakers who lack a capacity to tell beeches and elms apart, the words “beech” and “elm” are synonymous? That seems implausible. Putnam suggests that, when it comes to such *natural kind* terms, we rely on a division of linguistic labor. (Natural kind terms – “gold,” “water,” “planet,” “tiger” – designate stuffs and objects thought to occur naturally, and are distinguished from *artifactual kind* terms: “table,” “senator,” “dollar bill.”) We use “elm” and “beech,” for instance, to designate species of tree *that would be so labeled by experts*.

Twin Earth cases and the phenomenon of the division of linguistic labor make it clear that agents who are indiscernible in all relevant intrinsic respects could nevertheless differ in what they mean by their utterances. If this is right, accounts of meaning that focus solely on agents considered in isolation are bound to fail. An adequate account of meaning apparently brings with it a battery of social and contextual elements. Putnam puts it succinctly: “Cut the pie any way you like, ‘meanings’ just ain’t in the *head!*” (1975b: 227). The philosophical impact of this thesis – I shall call it *externalism* – would be hard to overstate. As we shall see, what goes for meaning goes for thought as well. If Putnam is right, the traditional conception of the mind as a spectator on the “external world” must be abandoned, replaced by a conception of the mind as constituting – and constituted by – the world.

This is to get ahead of our story, however. Let us look first at Putnam’s articulation and defense of functionalism, a conception of the mind according to which minds comprise systems of relations among elements that resemble the states of a computing machine.

Philosophy of mind

In the 1950s, English-speaking philosophers under the spell of Wittgenstein came to regard philosophical questions as expressions of linguistic befuddlement. We ask, for instance, “What is truth?” and interpret this as a substantive question, one that calls for the investigation of some independently existing reality. Wittgenstein argued that such questions occur to us only when we distance ourselves from linguistic practices that give form to talk of truth (or any other philosophically challenging concept).

Augustine (*Confessions*, XI, xiv) remarked about time: “What then is time? If no one asks of me I know; if I wish to explain to him who asks, I know not.” The sentiment (though not Augustine’s subsequent treatment of it) is profoundly Wittgensteinian. So long as we *use* language in the pursuit of ordinary human ends, we remain innocent of philosophy. We are moved to philosophical questioning when we step outside the linguistic practices that ground our use of words. We ask “What is time?” and seek an answer in a way that ignores the way “time” and its cognates are actually deployed in our linguistic community. Once we lose our moorings within language, our theorizing is colored by a misapprehension of

the roles of terms that generate familiar philosophical puzzles. This misapprehension is systematic: the same kinds of theory arise over and over in the history of philosophy.

Wittgenstein's positive proposal is deflationary. Philosophical puzzlement requires treatment. When a philosopher re-immerses himself in the linguistic practices and forms of life that give sense to the terms he finds bewildering, the bewilderment ebbs. Philosophical questions are not answered but laid to rest. The temptation to pose such questions, although perfectly natural, requires a kind of therapy (see WITTGENSTEIN). The philosopher who responds to this therapy is no longer impelled to philosophize; the end of philosophy is the end of philosophy.

Wittgenstein's approach to philosophical issues concerning the mind led to the rejection of the traditional idea of minds as mental organs that receive inputs via the senses and yield outputs in the form of utterances and bodily motions. "Mind" is a substantive noun, but talk of minds is not talk of a substance or entity associated with, but somehow distinct from, the body. On the contrary, in regarding you as possessing a mind, I regard you as engaging in intelligent activities, responding to the world in intelligible ways, and so on. Thoughts like these led, in turn, to philosophical behaviorism (see RYLE): possessing a mind is exclusively a matter of behaving, or being disposed to behave, in particular ways (see, e.g., Ryle 1949; for a response, see Putnam 1975b, chs 14, 15, 16). Behaviorists hoped to analyze or translate talk of mental goings-on (feelings, thoughts, intentions) into talk of behavior and behavioral dispositions. To be depressed, for instance, is not to be in a particular kind of inner state, but to mope about, complain, or be disposed to complain, and the like. Behaviorists need not deny that inner states *accompany* bouts of depression, only that these inner states *are* the depression.

One difficulty for the behaviorist program stemmed from the fact that behaviorist analyses of mental concepts typically included reference to other mental concepts. Someone who is depressed, for instance, is disposed to form thoughts of certain sorts, and to acquire (or lose) certain motives and desires. When we attempt to analyze *these* mental concepts behavioristically, we find we must appeal to other mental concepts; analyses of these concepts require reference to still other mental concepts; *and so on*. (As we shall see, the interconnectedness of mental concepts comes to the fore with the development of behaviorism's intellectual successor, functionalism.)

The analytical program of behaviorism was challenged, first by the advent of the mind-brain identity theory (see Place 1956, Smart 1959) and then by functionalism. Mind-brain identity theorists defended the thesis that conscious states were at bottom states of brains. They argued that the kinds of correlation known to hold between subjects' reports of states of consciousness and states of those subjects' brains are best construed as evidence for the identification of states of consciousness with brain states. Imagine that, while shopping, you drop a can of tomato soup on your foot and, as a result, you experience a throbbing pain in your big toe. Neuroscientists tell us that, when you experience a pain of this sort, certain kinds of event occur in your brain. (Let us pretend that pains are associated with the firing of C-fibers in the spinal cord.) Identity theorists argued that the best explanation of the correlation between C-fiber firings and reports of pain was that being in pain just is the firing of C-fibers.

Functionalism

Nowadays many scientifically minded theorists regard it as close to obvious that states of mind are brain states, mental events are neurological events. Professional philosophers, however, have by and large resisted this conclusion. This is not because philosophers have a preference for mind–body dualism, but because most philosophers have been convinced by arguments pioneered by Putnam that the mind–brain identity theory suffers a fundamental defect (see Putnam 1975b, chs 18, 19, 20).

Consider the fact that we unhesitatingly ascribe states of mind to creatures other than human beings. Think of being in pain, and suppose for the sake of argument the mind–brain identity theory were correct: pains *are* brain states; your being in pain *is* your being in a particular kind of brain state: your C-fibers are firing. So far, so good. But now consider: can an octopus feel pain? It surely seems so. The neurological makeup of an octopus is very different from the neurological makeup of a human being, however. This seems to imply that octopodes, sporting a different physiology, lack a capacity for pain: if pains are C-fiber firings, and octopodes' pain responses are triggered by different mechanisms (as they surely are), then octopodes do not feel pain! Suppose we encountered intelligent creatures from distant planets who were like us in many ways but whose biology was silicon-based. We might have excellent grounds for regarding such creatures as experiencing pain, yet, if the mind–brain identity theory were true, this would be impossible if such creatures lacked C-fibers (as they almost certainly would). If having a pain is a matter of being in a particular kind of neural state, no creature lacking such states could experience pain.

If human beings, octopodes, and Alpha Centaurians can all experience pain, then it is hard to see how pain could be identified with kinds of brain state found only in human beings (and their near relations). Suppose, however, we think of pain states, not as neurological states, but on the model of *computational* states. Reflect on an ordinary desktop computer. The device's operation is governed by *programs* that it runs. When you elect to print a document you have been working on, your desktop computer runs a simple program that sends signals to a printer, which then prints the document. Now suppose we distinguish between the program your desktop computer is running and a particular physical *implementation* of that program. The machine's running the program is a matter of its going into a sequence of physical states. These states, we might say, *realize* the program. Note, however, that a different machine could run the very same program by going into a sequence of very different kinds of physical state. In the 1950s, computing machines consisted of ungainly arrays of vacuum tubes; modern computers make use of tiny transistors; in the nineteenth century, Charles Babbage constructed a sophisticated computing machine using brass gears and cylinders; and today there is talk of molecular computers. It is possible for all of these devices to run the very same program, to engage in the very same sequence of computations, and so to encompass the very same computational states.

Distinct machines can be in the same computational state, then, even if they are made of very different physical ingredients. All that is required is an isomorphism – a one–one correspondence – between sequences of operations performed by the machines (and sameness of inputs and outputs). You feed into a simple calculator “7,” “+,” and “5,” and the calculator displays “12.” The causal chain leading the calculator through

this computation has a certain physical character. When you type this same sequence into your desktop computer or dial it into a Babbage machine, these devices go through vastly different kinds of causal sequence to arrive at the same output: "12." At any rate, the sequences are vastly different considered solely as physical events. They exhibit, however, a common structure, a corresponding set of relations. You might put this by saying that, considered concretely, the events are very different but, considered at a higher level of abstraction, the sequences they embody are the same.

What has any of this to do with the mind? In suggesting that states of mind are computational states, Putnam is not imagining that creatures with minds – human beings, for instance – are "mere robots," creatures whose actions are inflexible and "mindless." The idea, rather, is that states of mind owe their identity, not to their physical makeup, but to their place within a structured system. To avoid misleading associations, I shall speak henceforth, not of computational states, but of functional states. (Are functional states and computational states co-extensive? This is controversial. A computational state can be given a particular sort of formal characterization. If every functional state or process is characterizable in this formal way, then every functional state is a computational state.)

Functional states are picked out by reference to roles they occupy – functions they perform – within a system. An analogy may help. Wayne is a vice-president of the Gargantuan Corporation. What exactly does Wayne's being a vice-president amount to? Wayne is 175 cm tall, balding, and overweight. These intrinsic properties of Wayne seem not to bear on his being a vice-president. Wayne could be "re-orged," and replaced by Becky, a petit brunette, by Oscar, a robot, by Hans, a chimpanzee fluent in sign language, or even by Renée, an immaterial angel. Wayne is a vice-president, not in virtue of his intrinsic properties, but in virtue of relations he bears to others in the organization in which he occupies this office. Anyone (or *anything!*) bearing these relations would be a vice-president.

Functionalists contend that what goes for vice-presidents goes for states of mind. Being in pain is not a matter of being in a particular kind of neurological state, but a matter of being in a state that bears the right sorts of relation to other components of the system to which it belongs. In your case, a particular neurological state occupies the pain role; in the case of an octopus or an Alpha Centaurian, very different kinds of physical state fill the pain role. A creature is in a state of pain when it is in a state that is typically caused by tissue damage, and that causes certain characteristic beliefs and desires (the belief that this hurts, for instance, and a desire for the pain to stop), and certain characteristic actions (if you've stepped on a tack you will quickly move your foot). This makes functionalism sound like dressed up behaviorism. Functionalism, however, unlike behaviorism, does not require that states of mind be characterizable solely in terms of stimuli and responses. How you respond to pain – your behavior – can depend partly on what you believe and desire. If you are trying to impress a companion with your toughness, you may shrug off a pain that you would react to very differently were you alone.

You may worry that this way of characterizing mental states is ultimately circular. We designate a mental state by noting its relations to other mental states. These, in turn, are characterized by reference to other mental states. Eventually we come back to the original states.

The threat of circularity is warded off by means of a technique introduced in a different context by Frank Ramsey and refined by David Lewis (see Lewis 1972). The issues are technical, but the guiding idea is straightforward. Imagine that you define states of mind by locating them as nodes in a network of nodes, each of which represents a distinctive kind of mental state. The system of nodes is anchored at one end by relations to incoming stimuli, and at the other end by behavioral outputs. Now the pain node will have a certain unique structural relation to other nodes in the system; a feeling of pleasure will have another kind of structural relation; and a belief or desire will exhibit other kinds of structural relation. We can then say that being in pain is a matter of being in a state exhibiting *these* kinds of relation to elements in a system with *this* kind of structure.

Putnam summarizes this line of reasoning by describing states of mind (indeed computational or functional states generally) as *multiply realizable*. This means that the very same state of mind can be realized by many different kinds of physical (or perhaps non-physical, ectoplasmic or angelic) state. If one state can have many realizers, that state cannot be identified with or reduced to any of its realizers. Thus, although states of mind are possessed by ordinary conscious agents *by virtue of* those agents' being in some physical realizing state, mental states are not reducible to the physical states that realize them – or so functionalists contend.

Despite its immense popularity, functionalism has been widely criticized. Putnam himself has been among its most vocal critics. Even so, it is fair to say that functionalism remains hugely influential, both inside and outside philosophy. Functionalism provides a way of understanding how mentality could be housed in the brains of human beings (and in the nervous systems of other intelligent species). In addition, functionalism leaves room for distinctive *levels* of explanation. We might come to understand the behavior of a computing machine by investigating its physical makeup or by studying its program. In the same way, you might explain my behavior by citing complex processes in my central nervous system or by reference to my beliefs, desires, and intentions. Of course, the physical makeup of a computing machine might be extremely complicated, and the physical makeup of a human being more complicated still. In most cases, a purely physical explanation of the behavior of either would be a practical impossibility. Nevertheless, functionalism provides a way of seeing how we could be warranted in offering “higher-level,” functional explanations of the behavior of complex systems, and doing so in a way that does not compete with lower-level, purely physical, explanations.

Functionalism spurned

Recall Putnam's line on meaning: the meaning of your utterances depends, not merely on your intrinsic features, but on relations you bear to your surroundings. When you utter “that's water,” for instance, your utterance concerns water (H₂O) only if you stand in an appropriate relation to water. When your Twin Earth counterpart produces an indistinguishable utterance, that counterpart says something different. When you speak of elms and beeches what you mean is determined in part by experts in your linguistic community who are in a position to identify and distinguish elms and beeches. In this regard, meanings are community affairs. You and the expert mean the same

when you speak of elms, even though you lack the expert's knowledge of the distinguishing marks of elms. Your beliefs about elms might be largely false, still you mean by "elm" what others in your linguistic community mean: your talk of elms is talk of *elms*.

So far, this is a thesis about the meaning of utterances. The thesis is easily extended, however, to the contents of our thoughts: what those thoughts concern. Imagine that, on Earth, Debbie is anticipating a cool drink of water on a hot day. Debbie entertains a thought she would express by saying "that's water." At the same time Debbie's counterpart on Twin Earth, Twin Debbie, is entertaining a thought *she* would express by saying "that's water." Debbie's utterance and thought concern water. Twin Debbie's utterance and thought, in contrast, are not about water; water, after all, is H₂O, and Debbie's utterance and thought are not about H₂O; they are about XYZ, *twin* water! Of course, Twin Debbie *calls* twin water "water," but that is another matter. Debbie and Twin Debbie's uses of "water" resemble the use of "burro" by a Spanish speaker and an Italian. In the mouth of a Spanish speaker, "burro" means donkey; uttered by an Italian, "burro" means butter.

Putnam holds that cases like these make it clear that what our thoughts concern, as well as what our words mean, is fixed, not solely by what is inside our heads but by relations we bear to the world around us. "Water" in Debbie's (but not Twin Debbie's) mouth means water in part because Debbie (but not Twin Debbie) stands in an appropriate causal relation to water – H₂O. Similarly, thoughts Debbie (but not Twin Debbie) would express by utterances featuring the word "water" concern water – H₂O – in part because Debbie (but not Twin Debbie) stands in an appropriate causal relation to water. Twin Debbie's "water" utterances and thoughts she would express by means of these utterances concern, not water, but twin water, XYZ. To be sure, there is no relevant *internal* difference between Debbie and Twin Debbie (ignoring the fact that Debbie's constitution includes H₂O, Twin Debbie's, XYZ). The contents of our thoughts, like the meanings of our words, depend on our context, most particularly on causal relations we bear to objects in the world and social relations we bear to others in our linguistic community. These external relations are partly constitutive of the meanings of words and the contents of thoughts.

It is natural to extend the externalist thesis that meanings are not "in the head" to the meanings or contents of states of mind. This seems to imply that states of mind, or at any rate the contents of states of mind, are not in the head! (If you think that the content of a state of mind is essential to it – the belief that snow is white is essentially the belief that snow is white – then an externalism about content straightforwardly yields an externalism about states of mind with content: beliefs, desires, intentions, and the like.) Surely this is ridiculous! Or is it?

Before trying to answer this question, let us reflect on the implications of all this for functionalism. A functional state is a state of a system, a state definable wholly by relations it bears to other states of the system (and to inputs and outputs). Twin Earth cases, however, appear to show that distinct agents (Debbie and Twin Debbie, for instance) could be functionally identical yet differ mentally: one is thinking of water, another of twin water. You might regard functionalism as framing a more or less traditional "internalist" conception of the mind and its contents. If you are attracted to externalism – a broadly contextual account of the

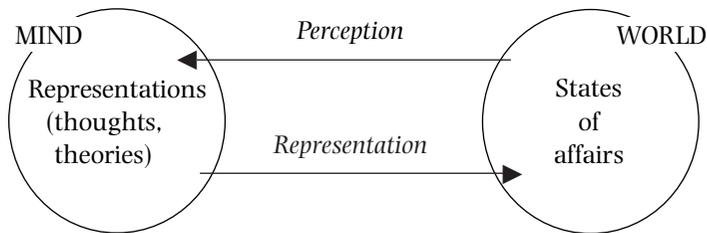
contents of states of mind – you thereby abandon the traditional view of minds as self-contained entities radiating thoughts onto an “external world.” The boundary between mind and world becomes blurred: “the mind and the world jointly make up the mind and the world” (Putnam 1981: xi).

If functional states are purely internal states of a system (states definable by their relations to other states of the system), however, and if states of mind are not, then states of mind are not functional states. We are in this way led back to the idea that making sense of the mind and its contents is a matter of seeing intelligent agents in context. The attempt to locate minds inside heads is analogous to an attempt to characterize chess pieces solely by reference to their intrinsic features.

Metaphysics

All this brings us to Putnam’s attack on “metaphysical realism,” a doctrine, rarely articulated but widely taken for granted, according to which the mind and the world are separated by an epistemological chasm. Minds respond perceptually to, and represent how things stand in, the world. Nevertheless minds are depicted as occupying a standpoint *outside* the world.

Although most of us are by and large committed to the view that minds are physical constituents of the world, it can still seem perfectly natural to *represent* the relation minds bear to the world in a way that situates thoughts “in here” and their worldly objects “out there.” We are spectators on the world. Perceiving is a matter of the world affecting the mind (the incoming arrow in the figure), and thinking about the world is a matter of aiming thoughts at the world (the outgoing arrow).



Descartes’s mind–body dualism is only one extreme form of metaphysical realism. Indeed, Putnam holds that the metaphysical force of Cartesian dualism lies not in the contention that minds are immaterial entities, but in the idea that minds stand apart from the world on which their thoughts are directed. This picture, the core of metaphysical realism, survives the transition to materialist conceptions of the mind. The result is a kind of internal instability. On the one hand, modern science encourages us to regard minds as material objects alongside other material objects. On the other hand, metaphysical realism depicts the mind as a spectator on the world, seemingly locating minds outside the world they represent. If we insist on situating minds in the world, however, we must abandon metaphysical realism.

One way to get at all this is to consider the implications of meaning externalism. Externalism undermines the conviction that meanings and mental contents are “in the head.” What you mean when you utter an English sentence and what your thoughts are directed on, depends in part on your circumstances (and not merely on your constitution or internal organization). The point can be illustrated by means of a simple analogy (borrowed from Wittgenstein). Imagine a picture of a smiling face. Imagine the face appearing in the foreground of a depiction of a child’s birthday party. Here the face expresses benevolent happiness. Now imagine the face situated against a background of horrible suffering. In this context, the face expresses evil. Perhaps our thoughts are like this. The same *form* of thought in one context expresses one content, and, in a different context expresses something entirely different. (“Form” here refers to the “shape” or “intrinsic character” of a thought. “Burro” in Spanish encompasses donkeys, in Italian it designates butter. The same linguistic form can have different meanings in different linguistic settings.)

This is one component of Putnam’s attack on metaphysical realism. A second component is epistemological. In sharply distinguishing representations of the world – beliefs and theories – and the world those representations concern, we create an unbridgeable chasm. Our representations purport to “match” reality, but we are in no position ever to effect a comparison. At best we can measure representations against other representations. Suppose you believe the ice is thin. You decide to check your belief by examining the ice. You are not measuring your belief against the ice, but measuring your belief against other perceptually induced beliefs about the ice.

If you regard this as a natural and unavoidable feature of the human predicament, you are at least a closet metaphysical realist. One consequence of such a view is that it opens the door to “external world skepticism”: what grounds could we have for thinking that our beliefs about the “external world” are true? If we have access only to our own representations, then we could have no assurance that those representations “match” the reality they purport to represent, or even whether there is any external reality beyond the representations themselves!

The situation is one Descartes dramatized by imagining an evil demon. The evil demon has the power to make our beliefs about the world false. Descartes’s attempt to reconcile metaphysical realism with the conviction that, properly pursued, knowledge of the external world was attainable, involved appeal to a benevolent God: God is such that he would not let us err concerning truths we find indubitable. This gives Descartes a foundation on which to erect an account of knowledge according to which we are entitled to be confident that our beliefs are true provided those beliefs are not based on unproven assumptions. The argument’s appeal to a benevolent God, however, strikes most readers as unconvincing.

Note that the skeptical challenge – what entitles us to suppose that our representations of the world match the world? – presupposes metaphysical realism. Skepticism is metaphysical realism seen in the mirror of epistemology (Heil 1998). Metaphysical realism makes the skeptical question inevitable, and the skeptical question makes sense only given the kind of mind–world separation that makes up the core of metaphysical realism. A refutation of realism, then, can be seen as a refutation of “external world” skepticism – and vice versa. This is precisely Putnam’s strategy.

Brains in vats

Appeals to a benevolent God aside, let us follow Putnam in updating the skeptical challenge by asking what grounds we have for believing that we are not *brains in vats*. Imagine that you have been kidnapped by an evil scientist, drugged, and your brain removed from your body and kept alive in a vat of nutrients. Nerve endings previously attached to bodily organs are attached now to a super computer. The computer precisely simulates incoming nerve impulses. Nervous stimulation that two days ago would have come from your retina, for instance, now issues from the computer. As far as you can tell, the world is unchanged. Your visual and auditory experiences, even the kinesthetic feedback you receive when you seem to move your body are fed to you by the computer. All this, although undoubtedly fanciful, seems at least physically possible. But, the skeptic insists, if it is *possible*, how could we ever be in a position to know (or even reasonably believe) that we are not brains in vats?

Note, first, that the brain-in-a-vat possibility is a possibility only so long as we accept metaphysical realism and its attendant gap between how the world is and how we represent it as being. Does this give us a reason to abandon realism? It hardly seems so. If realism implies that we might be brains in vats, then this is something we shall have to live with. (If you are inclined to dismiss the possibility as idle, ask yourself what grounds you have for dismissing it.)

Here Putnam goes on the offensive. Suppose metaphysical realism implies that we might be brains in vats. If we could establish that it is *not* possible that we are brains in vats, then we will have established that metaphysical realism is false. But how could anyone hope to prove that it is not possible that we are brains in vats? We have, after all, granted that the envisaged envatting of a brain lies within the realm of physical possibility.

Recall Putnam's take on meaning. The meaning of an utterance (or the content of a thought you might express with that utterance) depends on context, most especially on relations speakers and thinkers bear to their surroundings. In the simplest case, your thought of a tree concerns this tree because this tree (and no other) is causally responsible for it. Suppose we generalize this observation. The English words "brain" and "vat" mean what they do in part because members of the English-speaking linguistic community stand in appropriate causal relations to brains and vats. Similarly, your thoughts of brains and vats, thoughts you might express using the terms "brain" and "vat," concern brains and vats because you are, as an English speaker, an agent standing in appropriate causal relations to brains and vats. The causal relations in question are no doubt complex, and it would be difficult to spell them out in detail. But, as the Twin Earth cases seem to show, the presence of an appropriate causal connection is at least a necessary condition for our words and thoughts connecting to the world. Debbie's utterance of "water" designates water, and thoughts she would express using this term concern water, in part because Debbie's is causally related to water. Twin Debbie's utterances of "water" differ in meaning and her corresponding thoughts differ in their content because Twin Debbie stands in comparable causal relations, not to water, but to XYZ, twin water.

Let us allow that "vat" and "brain" mean what they do in part because those of us who deploy these terms stand in certain causal relations to vats and brains. Now

consider an envatted brain, Evan. Consider, in particular, Evan's causal links to the world outside the vat and their bearing on the significance of Evan's "utterances" of "brain" and "vat" (and thoughts Evan might "express" using these terms). Sources of stimulation for these utterances and thoughts are not brains and vats, but electrical events inside the super computer to which Evan is attached. These electrical events – stand-ins for real brains and vats – produce sensory experiences in Evan that precisely resemble the sensory experiences you might have when you encounter a vat or a brain. Evan's situation obliges us to reconstrue Evan's utterances and thoughts, just as we did in the Twin Earth case. When Evan has a thought that he might express by uttering the sentence "That's a brain," we should have to interpret his utterance as meaning something like "That's electrical state s_1 ." Similarly, when Evan harbors a thought he would express by uttering "That's a vat," we must interpret this utterance as expressing something like what we would express in English as "That's electrical state s_2 ."

Evan's utterances will need to be systematically reinterpreted. His utterances *formally* resemble English utterances, but they differ significantly in what they *mean*. We can mark this systematic difference by describing Evan as "speaking," not English, but Vat-English, just as your twin on Twin Earth speaks Twin English, not English. Of course, just as your twin calls Twin English "English," so Evan calls Vat-English "English." When Evan "says" "I speak English," his "utterance," translated into English (*our* language), means "I speak Vat-English," and this utterance is true.

With all this as background, we are in a position to appreciate Putnam's anti-skeptical argument. Suppose you entertain a thought that you would express by uttering the sentence "I am a brain in a vat." If you are an English speaker, then this sentence is false. Why? If you are an English speaker, you are connected in a normal way with brains, vats, trees, and the like and *not* plugged into a super computer. If you are an English speaker, then, you are not a brain in a vat.

Astute readers will be quick to point out that it is hard to take much comfort from this fact. True, if we grant meaning externalism, we are not brains in vats *if we are English speakers*. But what gives us the right to assume that we *are* English speakers? After all, if we were brains in vats, we would not be English speakers! It looks as though, in order to know that we speak English (and not Vat-English) we should first have to know that we are not brains in vats. We cannot, then, under pain of circularity, appeal to this fact (or alleged fact!) to establish that we are not brains in vats.

Let us think a little harder about this. Consider Evan, and his "utterance" of "I am a brain in a vat." Evan, as we know, is envatted and "speaks," not English, but Vat-English. Evan's "utterance" of "I am a brain in a vat" translated from Vat-English into English would mean something like "I am a computer state of type s_n ." But this utterance is manifestly false: Evan is a brain in a vat, not a computer state! It appears that you can generalize this point and use it in simple argument to the conclusion that you are not a brain in a vat:

- 1 If I am a brain in a vat, I express a falsehood in uttering the sentence "I am a brain in a vat."
- 2 If I am not a brain in a vat, I express a falsehood in uttering the sentence "I am a brain in a vat."
- 3 I am a brain in a vat or I am not a brain in a vat.

- 4 In uttering the sentence "I am a brain in a vat," I express a falsehood. (From (1), (2), and (3))
- 5 In uttering the sentence "I am a brain in a vat," I utter a sentence meaning that I am a brain in a vat.
- 6 I am not a brain in a vat. (From (4) and (5))

This argument is valid: its premises logically imply the conclusion. But is the argument sound? Does it establish what it purports to establish? Does the argument enable you to prove that you are not a brain in a vat?

Before taking up this question, let us remind ourselves why the issue is important for Putnam. Metaphysical realism, in bifurcating mind and world, implies that it is possible that we are massively deluded: it is possible that we are brains in vats. If we can exclude this skeptical possibility, we shall have thereby established the inadequacy – the falsehood or perhaps incoherence – of metaphysical realism.

Does Putnam's argument work? Suppose that, after studying the argument you conclude that it fails to show that you are not a brain in a vat: *you might be a brain in a vat anyway*. Suppose you express this possibility via the sentence "I am a brain in a vat." As we have seen, your utterance of this sentence is false if you are an English speaker, and it is false if you speak Vat-English. In either case it is false. Assuming (for simplicity) these are the only possibilities, the sentence must be false. Yes, you might reply, but it is vital to know what the false sentence *means*. Premise (5) of the argument tells us that the sentence means that you are a brain in a vat. If you knew this, and knew as well, that the sentence was false, you *would* know that you are not a brain in a vat. But why should we accept premise (5)?

Remember, we are supposing that you are running through the argument in an effort to establish that you are not a brain in a vat. Imagine that you have just run through premise (5). How *could* premise (5) express a falsehood? To see the difficulty in denying (5), pretend that we are eavesdropping on Evan as he runs through the argument (the computer connected to Evan includes a loudspeaker so that we can eavesdrop on Evan's ruminations). When Evan reaches premise (5), he concludes: "In uttering the sentence 'I am a brain in a vat,' I utter a sentence meaning that I am a brain in a vat." Evan is "speaking" Vat-English, so we shall need to translate this utterance into English. When we do so we obtain something like: "In uttering the sentence 'I am a brain in a vat,' I utter a sentence meaning that I am a computer state of type s_n ." This utterance is certainly true. More generally, utterances of the form "in uttering the sentence '*P*', I utter a sentence meaning *P*," are bound to be true.

If Putnam is right, then the meanings of these sentences depend on all sorts of causal and contextual factors. We need know nothing of these, however, for them to determine the meaning of what we say and the content of what we think. Debbie's utterances of "water," and thoughts she would express by such utterances, concern water, not because Debbie has figured out that she is causally connected to H_2O (and not XYZ), but because Debbie is causally connected to H_2O (and not XYZ).

Imagine, again, that we are eavesdropping on Evan running through Putnam's argument. Pretend that Evan finds the argument convincing, concluding in an excited tone of voice: "I am not a brain in a vat!" Surely, you think, Evan is deluded. As we can plainly see, Evan *is* a brain in a vat. This is too quick, however. Before we can evaluate

Evan's conclusion, we must translate it from Vat-English into English. When we do so, we obtain something like: "I am not a computer state of type s_n !" This, of course, is true; Evan is *not* deluded. We can apply this lesson to our own case: given meaning externalism, we could not be deluded in concluding from Putnam's argument that we are not brains in vats.

Implications for metaphysical realism

Evan is certainly not deluded in the sense of believing a falsehood. He believes that he is not (as we should put it) a computer state, and he is correct. Indeed, most of Evan's beliefs about his actual situation are correct. (The argument is not affected by supposing that Evan has false beliefs; all of us have our share of false beliefs.) Of course Evan cannot appreciate that he is a brain in a vat hooked to a computer programmed by an evil scientist. The thought is not one Evan is in a position to entertain.

Perhaps this is all Putnam needs to show that metaphysical realism is untenable. Metaphysical realism presumes an epistemological gap between what we take to be the case and what is the case. One way to express this is in terms of the possibility that our beliefs about what is the case are massively false. If we accept externalism about the meanings of our words and the contents of our thoughts, if we suppose that what we mean and what our thoughts concern is fixed by our circumstances, we thereby exclude the possibility of massive error. This is what Putnam's argument shows.

If we return to Evan, however, we can see that, although Evan may not be massively deceived, a gap remains between what he takes to be the case and what is the case. The fact, if it is a fact, that Evan is in no position to entertain thoughts concerning what is the case – thoughts about brains and vats – provides scant comfort when we consider our own circumstances. We can still envision a gap, even if we cannot envision what might lie on the far side of the gap. And this evidently leaves realism standing.

Is this unfair? We are imagining that our circumstances might be wildly different from what we take them to be even though our beliefs about those circumstances are, on the whole correct: the deep truth is not thinkable by us. But what sense could be made of the suggestion that things might be some way, although we cannot so much as consider what that way might be? This sounds like nonsense; and a nonsensical possibility is no possibility at all.

Readers familiar with Berkeley will recognize this line of reasoning. Berkeley dismisses the possibility of a material world, a world of objects existing mind-independently, on the grounds that we cannot so much as entertain thoughts concerning such a world. If we cannot entertain thoughts concerning X , then plainly we could have no reason to think X exists or might exist: talk of X is empty.

The situation we have been envisaging, however, is not one in which we endeavor to think the unthinkable, but one in which we acknowledge our fallibility, recognizing that we could be wrong about almost anything without being in a position to entertain thoughts as to how things actually are. Perhaps this is all a realist needs: it is possible that reality (or a significant portion of reality) is not just unknown, but unknowable by us owing to our circumstances. Or perhaps a realism of this sort leaves behind the traditional impetus for realism. In either case, Putnam's reflections push realists – and

their bedfellows, the skeptics – to examine their fundamental assumptions. For that, realists and anti-realists alike should be grateful.

Ontological pluralism

We have been focusing on Putnam's contention that the world is not mind-independent: "the mind and the world jointly make up the mind and the world." But there is another dimension to Putnam's dissection of metaphysical realism (see Putnam 1987, lecture 1). A metaphysical realist regards the world as possessing a definite character quite independently of our ways of thinking about it. We represent this character in various ways: truly or falsely, subtly or clumsily. Our everyday beliefs about the world represent it as being one way, for instance; the sciences represent it differently. The "scientific image" and the everyday, "manifest image" of the world appear in various ways to be at odds. (Talk of scientific and manifest images originated with Wilfrid Sellars (see SELLARS); see Sellars 1963, ch. 1.) The surface of the desk at which I am sitting appears smooth and continuous. Physics, however, tells us that the desk is a cloud of particles, widely spaced and in constant motion. Which description of the desk is the correct one? Perhaps the apparent desk, the desk of the manifest image, is a *mere* appearance.

This is the reaction of the metaphysical realist. The realist starts with the idea that the world is a single definite way. We can describe the world in many different ways; some kinds of description capture the world better than others, however. My ordinary description of my desk, for instance, is at best a crude approximation. Taken literally it is false. We can edge closer to the truth by turning to physics. When we do, we learn that the world contains no desks, only clouds of invisible particles.

Can we make sense of this picture? Return to my desk. How many things are stacked on it? An answer to this question will depend on how we decide to count. We could, for instance, count pencils, pens, books, and memos. We could just as easily count pages of books, parts of pens and pencils. Or we could count particles of which all these things are composed. What is *the* correct way to count? What is *the* correct answer to the question, how many things are on my desk? Such questions are wrong-headed. There are *many* ways to sort objects on my desk, *many* correct answers to the original question. The contents of my desk can be "carved up" in different ways. How we do so depends on *us*: our aims or purposes.

We deploy systems of concepts in representing the world. The metaphysical realist sees these concepts as matching well or badly what is "out there." But this is the wrong model. The concepts we use determine, rather than merely reflect, what is out there, at least in the sense that they determine objects' boundaries, hence what is to count as an object. It is not that there are no divisions in nature, but that there are too many. Our concepts, or rather systems of concepts, make some of these divisions salient. In saying how the world is, we invoke one or another conceptual system or scheme. Which scheme we invoke depends in part on features of us, our needs and interests. As we learn more and as our needs and interests change, our conceptual schemes evolve.

Suppose this is right. It is then hard to see how we could make sense of talk of a world independent of any conceptual scheme (Kant's noumenal world, the "thing in itself"). We could have no way of describing that world, no way of thinking it: describ-

ing and thinking involve representing in terms of a conceptual scheme. Consider a map of the surface of the Earth. We can depict the Earth's surface by means of a Mercator projection, a Peterson projection, a spherical projection. Imagine someone dissatisfied with these insisting that the Earth be depicted using no projection at all! This is what the metaphysical realist demands for representations of reality in general: a representational system or conceptual scheme that is utterly transparent. But a transparent scheme is no scheme at all.

Once you accept that there can be no sense in talk of a scheme-independent world, you may be moved to ask how competing schemes could be evaluated. Again, this is the wrong question. Just as Mercator projections and Peterson projections do not compete, so our modern scientific scheme does not compete with our everyday conception of the world. Both are entirely satisfactory on their own terms, both provide perfectly adequate depictions of our world. An everyday description can be wrong; I may falsely believe that there is a desk in my office. But if this belief is false it is not because science tells us that there are no desks (only clouds of particles). Ontology – what there is – is relative to a conceptual scheme. Desks do not figure in the ontology of the physicist, but this does not mean that desks are mere appearances. Insofar as we find it useful, or unavoidable, to deploy a conceptual scheme in which desks have a role, the ontological legitimacy of desks is assured. The metaphysical realist hankers after a single ontology: *the* ontology of the world. Instead, Putnam insists, we should embrace ontological pluralism: what there is depends in part on schemes or systems of concepts we find it convenient to deploy.

Externalism again

In assessing Putnam's attack on metaphysical realism, we have been granting externalism, the view that what we mean and what our thoughts concern depends in part on our circumstances. Putnam's defense of externalism relies heavily on Twin Earth cases: we imagine agents who are intrinsically indiscernible and yet whose utterances, and thoughts those utterances express, differ in significance.

Part of the idea here is that the projective character of thoughts – what is often called their *intentionality* – is due, not to intrinsic features of those thoughts, but to matters external to thinkers. Another of Putnam's examples nicely illustrates the point. Suppose you form a mental image of a particular tree, one in a nearby park, for instance. Now imagine an Alpha Centaurian, Fred, who lives on a planet barren of vegetation and so knows nothing of trees. One day Fred spills some paints that purely by chance form a design that you would regard as perfectly realistic representation of the tree in the park. Later, reflecting on the spilled paint, Fred forms a mental image indistinguishable intrinsically from your tree image. Is Fred imagining the tree in the park? That seems unlikely. If Fred is imagining anything, he is imagining a design produced by spilled paint. What accounts for the difference? The images are intrinsically alike, so the difference must lie elsewhere. Perhaps the difference stems from your being causally related to the tree in the park, in a way Fred is not. Your image of a tree *projects* to that tree because you stand in an appropriate causal relation to the tree; Fred's image projects, not to the tree but to the spilled paint, because the spilled paint, not the tree, plays the required causal role.

A view of this sort reverses the metaphor of projection. Thought does not project from the “inside out,” but from “outside in.” Must we go along? Perhaps not. Perhaps we can account for the projective character of thought by reference to intrinsic features of agents; perhaps projection is “inside out.” How, then, could we accommodate Twin Earth cases? Suspend doubt for a moment and pretend that the projectivity of a thought is like the beam of a flashlight radiating outward. What the beam of the flashlight illuminates depends both on the nature of the beam, as determined by intrinsic features of the flashlight, and on what happens to be “out there” to be illuminated. Just as flashlights on Earth illuminate water (H₂O), and flashlights on Twin Earth illuminate twin water (XYZ), so thoughts on Earth project to water, and thoughts on Twin Earth project to twin water. The moral? Twin Earth cases do not establish that the projective character of thought is due to incoming causal chains, then, only that what thoughts designate depends in part on the circumstances of thinkers.

Granted it is silly to compare the projectivity of a thought to the beam of a flashlight. Nevertheless it may be possible to base an account of the projective aspect of thought on intrinsic features of agents. Agents and their states of mind possess dispositionality, and dispositions are inherently projective. Locke’s example of a lock and key illustrate the idea. The key is *for* locks of a certain sort, and not for others. This is so even if no such lock has been manufactured (or if the lock the key fits is destroyed). The key is disposed to open one lock, but not another: the key *projects* to one lock, but not to another. The key’s so projecting does not depend on the key’s having been in causal contact with the lock, but solely on intrinsic features of the key (and intrinsic features of the lock).

Imagine now that an agent’s states of mind incorporate fine-tuned dispositions (Martin and Heil 1998; Martin and Pfeiffer 1986). Your thoughts of trees, for instance, project to trees, not perhaps because they are caused by trees, but because they dispose you to interact in appropriate ways with trees. To be sure, intrinsically indiscernible thoughts might dispose an inhabitant of Twin Earth to interact with twin trees. This does not show that your thought’s projective character comes from the outside, however. The projective character of those thoughts might be “built in” even if objects those thoughts “illuminate” depends on what objects are available to be illuminated.

Consider another much-discussed case (Davidson 1987). Don is wading through a swamp in a thunderstorm. Suddenly, a bolt of lightning reduces Don to a pile of ashes and simultaneously reconstitutes a nearby tree stump into a “molecular duplicate” of Don. Suppose that the molecular duplicate, Swampman, functions just as Don did prior to his sudden demise, and in particular Swampman has thoughts, images, and memories intrinsically indiscernible from Don’s. Swampman has many false memories. He seems to remember his twelfth birthday party, but he is in fact only a few minutes old. What of Swampman’s other thoughts and images, however? Swampman entertains thoughts and forms images intrinsically indiscernible from Don’s thoughts and images of trees, stars, water, and so on. Ought we to say that these thoughts and images are empty of significance until Swampman comes into causal contact with trees, stars, and water? Only someone with a prior commitment to an “outside-in” conception of thought would say so. Swampman’s mental condition includes finely-tuned dispositions that undergird the projective aspect of his thoughts. Of course, what those thoughts project *to* depends in some measure on what is “out there.” This is the lesson

of Twin Earth. But this need not lead us to imagine that the projectivity of thought must be explained by incoming causal chains.

Putnam's significance

This entry provides only the briefest introduction to one corner of Putnam's philosophical work. I have expressed reservations concerning two themes that have proved especially influential in recent philosophy. These critical comments afford, at best, only hints as to where a reader might disagree with those doctrines. Putnam's work is wide-ranging, rich, and interconnected in a way that undercuts piecemeal criticism. In attacking one Putnam thesis, a critic risks assuming positions that Putnam elsewhere rejects. Perhaps I have said enough to make it clear that Putnam's work is deeply insightful, penetrating, and synoptic. Even when Putnam self-confessedly goes up a blind alley, it is worth following him for the sake of observing some topic in a new and revealing light. (Besides, Putnam's blind alleys are more interesting than the well-trodden paths of other philosophers.) Hilary Putnam is one of a handful of philosophers who have individually shaped the fundamental character of contemporary philosophy.

Bibliography

Works by Putnam

- 1971: *Philosophy of Logic*, New York: Harper and Row.
 1975a: *Mathematics, Matter, and Method*, Philosophical Papers vol. 1, Cambridge: Cambridge University Press.
 1975b: *Mind, Language, and Reality*, Philosophical Papers vol. 2, Cambridge: Cambridge University Press.
 1978: *Meaning and the Moral Sciences*, London: Routledge and Kegan Paul.
 1981: *Reason, Truth, and History*, Cambridge: Cambridge University Press.
 1983: *Realism and Reason*, Philosophical Papers vol. 3, Cambridge: Cambridge University Press.
 1987: *The Many Faces of Realism*, 1985 Paul Carus Lectures, La Salle, IL: Open Court.
 1988: *Representation and Reality*, Cambridge, MA: MIT Press.
 1990: *Realism with a Human Face*, ed. J. Conant, Cambridge, MA: Harvard University Press.
 1992: *Renewing Philosophy*, Cambridge, MA: Harvard University Press.
 1994: "Sense, Nonsense, and the Senses: An Inquiry into the Powers of the Human Mind," *Journal of Philosophy* 91, pp. 445–517.
 1995: *Pragmatism: An Open Question*, Oxford: Blackwell Publishers.

Works by other authors

- Bohoss, G. (ed.) (1990) *Meaning and Method: Essays in Honor of Hilary Putnam*, Cambridge: Cambridge University Press.
 Clark, P. and Hale, B. (eds.) (1994) *Reading Putnam*, Oxford: Blackwell Publishers.
 Davidson, D. (1987) "Knowing One's Own Mind," *Proceedings and Addresses of the American Philosophical Association* 60, pp. 441–58.
 Dennett, D. (ed.) (1987) *The Philosophical Lexicon*, Newark, DE: American Philosophical Association.

JOHN HEIL

- Emerson, R. W. (1940) "Self-reliance," in *The Complete Essays and Other Writings of Ralph Waldo Emerson*, ed. B. Atkinson, New York: The Modern Library.
- Heil, J. (1998) "Skepticism and Realism," *American Philosophical Quarterly* 35, pp. 57–72.
- Lewis, D. (1972) "Psychophysical and Theoretical Identifications," *Australasian Journal of Philosophy* 50, pp. 249–58.
- Martin, C. B. and Heil, J. (1998) "Rules and Powers," *Philosophical Perspectives* 12, pp. 283–312.
- Martin, C. B. and Pfeiffer, K. (1986) "Intentionality and the Non-psychological," *Philosophy and Phenomenological Research* 46, pp. 531–54.
- Place, U. T. (1956) "Is Consciousness a Brain Process?," *British Journal of Psychology* 47, pp. 44–50.
- Ryle, G. (1949) *The Concept of Mind*, London: Hutchinson.
- Sellars, W. (1963) *Science, Perception, and Reality*, London: Routledge and Kegan Paul.
- Smart, J. J. C. (1959) "Sensations and Brain Processes," *Philosophical Review* 68, pp. 141–56.
- Wittgenstein, L. (1953) *Philosophical Investigations*, trans. G. E. M. Anscombe, Oxford: Blackwell Publishers.