

## The Continuity of Levels of Nature

William G. Lycan

Contemporary Functionalism in the philosophy of mind began with a distinction between *role* and *occupant*. As we have seen, the seductive compar-

ison of people (or their brains) to computing machines drew our attention to the contrast between a machine's program (abstractly viewed)

and the particular stuff of which the machine happens to be physically made, that *realizes* the program. It is the former, not the latter, that interests us *vis-à-vis* the interpretation, explanation, prediction, and exploitation of the machine's "behavior"; people build computers to run programs, and use whatever physical materials will best lend themselves to that task.

The distinction between "program" and "realizing-stuff," or more familiarly "software" and "hardware," lend itself happily back to the philosophy of mind when Putnam and Fodor exposed the chauvinistic implications of the Identity Theory. What "*c*-fibers" and the like are doing could have been done - this role could have been performed - by some physiochemically different structure. And sure enough, if the same role were performed, the same functions realized, by silicon-instead of carbon-based neurochemistry, or if our individual neurons were replaced piecemeal by electronic prostheses that did the same jobs, then intuitively our mentality would remain unaffected. What matters is function, not functionality; program, not realizing-stuff; software, not hardware; role, not occupant. Thus the birth of Functionalism, and the distinction between "functional" and "structural" states or properties of an organism.

Functionalism is the only positive doctrine in all of philosophy that I am prepared (if not licensed) to kill for.<sup>1</sup> And I see the "role"/"occupant" distinction (some say obsessively) as fundamental to metaphysics. But I maintain that the *implementation* of that distinction in recent philosophy of mind is both wrong and pernicious. And my purpose in this chapter is to attack the dichotomies of "software"/"hardware," "function"/"structure" in their usual philosophical forms, and to exhibit some of the substantive confusions and correct some of the mistakes that have flowed from them.

### The Hierarchy

Very generally put, my objection is that "software"/"hardware" talk encourages the idea of a bipartite Nature, divided into two levels, roughly the physiochemical and the (supervenient) "functional" or higher-organizational - as against reality, which is a multiple *hierarchy* of levels of nature, each level marked by nexus of nomic generalizations and supervenient on all those levels below it on the continuum.<sup>2</sup> See Nature as hier-

archically organized in this way, and the "function"/"structure" distinction *goes relative*: something is a role as opposed to an occupant, a functional state as opposed to a realizer, or vice versa, only *modulo* a designated level of nature. Let me illustrate.

Physiology and microphysiology abound with examples: *Cells* - to take a rather conspicuously functional term(!) - are constituted of cooperating teams of smaller items including membrane, nucleus, mitochondria, and the like: these items are themselves *systems* of yet smaller, still cooperating constituents. For that matter, still lower levels of nature are numerous and markedly distinct: the chemical, the molecular, the atomic, the (traditional) subatomic, the microphysical. Levels are nexus of interesting lawlike generalizations, and are individuated according to the types of generalizations involved. But cells, to look back upward along the hierarchy, are grouped into tissues, which combine to form organs, which group themselves into organ systems, which cooperate - marvelously - to comprise whole *organisms* such as human beings. Organisms, for that matter, collect themselves into organized (*organ-ized*) groups. And there is no clear difference of kind between what we ordinarily think of as single organisms and groups of organisms that function corporately in a markedly singleminded way - "group organisms" themselves, we might say.<sup>3</sup>

Corresponding to this bottom-up aggregative picture of the hierarchical organization of Nature is the familiar top-down explanatory strategy.<sup>4</sup> If we want to know how wastes and toxins are eliminated from the bodies of humans, we look for and find an *excretory system* interlocked with the digestive and circulatory systems. If we look at that system closely we find (not surprisingly) that it treats water-soluble and nonsoluble wastes differently. We find in particular a *kidney*, which works on soluble wastes in particular. If we probe the details further, proceeding downward through the hierarchy of levels, we find the kidney divided into renal cortex (a filter) and medulla (a collector). The cortex is composed mainly of nephrons. Each nephron has a glomerulus accessed by an afferent arteriole, and a contractile muscular cuff to control pressure (the pressure pushes water and solutes through the capillary walls into Bowman's Capsule, leaving blood cells and the larger blood proteins stuck behind). Reabsorption and so on are explained in cellular terms, e.g., by the special properties of the epithelial cells that line the

nephron's long tubule; those special properties are in turn explained in terms of the physical chemistry of the cell membranes.

The brain is no exception to this hierarchical picture of the organism and its organs. *Neurons* are cells, comprised of *somata* containing a nucleus and protoplasm, and fibers attached to those somata, which fibers have rather dramatically isolable functions; and we are told even of smaller functional items such as the ionic pumps, which maintain high potassium concentration inside. Neurons themselves are grouped into nerve nets and other structures, such as columnar formations, which in turn combine to form larger, more clearly functional (though not so obviously modular) parts of the brain. The auditory system is a fair example. There is evidence that the auditory cortex displays two-dimensional columnar organization:<sup>5</sup> columns of variously specialized cells arranged along one axis respond selectively to frequencies indicated by incoming impulses from the auditory nerve, while columns roughly orthogonal to these somehow coordinate input from the one ear with input from the other. The particular sensitivities of the specialized cells is to be explained in turn by reference to ion transfer across cell membranes, and so on down. For its own part, the auditory cortex interacts with other higher-level agencies - the thalamus, the superior colliculus, and other cortical areas - which interactions are highly structured.

Thus do an aggregative ontology and a top-down epistemology of nature collaborate. The collaboration has been eloquently argued for the science of psychology in particular, by Attneave (1960), Fodor (this volume), and Dennett (1978). I shall develop the point at some length, following Lycan (1981a).

### Homuncular Functionalism

Dennett (1978, p. 80) takes his cue from the methodology of certain AI research projects.<sup>6</sup>

The AI researcher *starts* with an intentionally characterized problem (e.g., how can I get a computer to *understand* questions of English?), breaks it down into subproblems that are also intentionally characterized (e.g., how do I get the computer to *recognize* questions, *distinguish* subjects from predicates, *ignore* irrelevant parsings?) and then breaks these problems

down still further until finally he reaches problem or task descriptions that are obviously mechanistic.

Dennett extrapolates this methodological passage to the case of human psychology, and I take it to suggest that we view a *person* as a corporate entity that corporately performs many immensely complex functions - functions of the sort usually called mental or psychological. A psychologist who adopts Fodor's and Dennett's AI-inspired methodology will describe this person by means of a flow chart, which depicts the person's immediately subpersonal agencies and their many and various routes of access to each other that enable them to cooperate in carrying out the purposes of the containing "institution" or organism that that person is. Each of the immediately subpersonal agencies, represented by a "black box" on the original flow chart, is in turn describable by its own flow chart, that breaks *it* into further, sub-subpersonal agencies that cooperate to fulfill *its* purposes, and so on. On this view, the psychological capacities of a person and the various administrative units of a corporate organization stand in functional hierarchies of just the same type and in just the same sense.

To characterize the psychologists' quest in the way I have is to see them as first noting some intentionally or otherwise psychologically characterized abilities of the human subject at the level of data or phenomena, and positing - as theoretical entities - the homunculi or subpersonal agencies that are needed to explain the subject's having those abilities. Then the psychologists posit further, smaller homunculi in order to explain the previously posited molar behavior of the original homunculi, etc., etc. It is this feature of the Attneave/Fodor/Dennett model that ingeniously blocks the standard Rylean infinite-regress objection to homuncular theories in psychology.<sup>7</sup> We explain the successful activity of one homunculus, not by idly positing a second homunculus within it that successfully performs that activity, but by positing a *team* consisting of several smaller, individually less talented and more specialized homunculi - and detailing the ways in which the team members cooperate in order to produce their joint or corporate output.

Cognitive and perceptual psychologists have a reasonably good idea of the sorts of subpersonal agencies that will have to be assumed to be functioning within a human being in order for that

human being to be able to perform the actions and other functions that it performs. Dennett (1978, ch. 9) mentions, at the immediately subpersonal level, a "print-out component" or speech center,<sup>8</sup> a "higher executive or Control component," a "short-term memory store or buffer memory," a "perceptual analysis component," and a "problem-solving component." And Dennett (1978, ch. 11) examines, in some clinical detail, a multi-leveled subpersonal structure that models the behavior that manifests human pain. "Behavior" here must be understood very richly, since Dennett scrupulously takes into account, not just the usual sorts of behavior that are common coin among philosophical Behaviorists and the apostles of commonsense psychology, but subtler phenomena as well: very small differences in our phenomenological descriptions of pain; infrequently remarked phenomena such as the felt time lag between our feeling that we have been burned and our feeling the deep pain of the burn; and (most interesting from the Homuncionalist point of view) the grandly varied effects of a number of different kinds of anesthetics and other drugs on a patient's live and retrospective reports concerning pain. Considerations of these various sorts serve the psychologists (and Dennett) as vivid pointers toward complexities in the relevant functional organization of the CNS, indicating the distinct black-box components at various levels of institutional organization that we must represent in our hierarchically arranged flow diagrams – the kinds of receptors, inhibitors, filters, damping mechanisms, triggers, and so on that we must posit – and the comparably various sorts of pathways that connect these components with each other and with the grosser functional components of their owners such as perceptual analyzers, information stores, and the speech center.

The homuncular approach, teleologically interpreted, has many advantages. I shall recount them when I have said a bit more about teleology. In the meantime, I put my cards on the table as regards the general form of a type-identification of the mental with the not-so-obviously mental: I propose to type-identify a mental state with the property of having such-and-such an institutionally characterized state of affairs obtaining in one (or more) of one's appropriate homuncular departments or subagencies. (The subagencies are those that would be depicted in the flow charts associated with their owners at various levels of institutional abstraction.) The same holds for mental

events, processes, and properties. To be in pain of type  $T$ , we might say, is for one's sub-personal  $\phi$ -er to be in a characteristic state  $S_T(\phi)$ , or for a characteristic activity  $A_T(\phi)$  to be going on in one's  $\phi$ -er.

### Homunculi and Teleology

It may be protested that the characterization " $\phi$ -er" and " $S_T(\phi)$ " are themselves only implicitly defined by a teleological map of the organism, and that explications of them in turn would contain ultimately ineliminable references to other teleologically characterized agencies and states of the organism. This is plausible, but relatively harmless. Our job as philosophers of mind was to explicate the mental in a reductive (and noncircular) way, and this I am doing, by reducing mental characterization to homuncular institutional ones, which are teleological characterizations at various levels of functional abstraction. I am not additionally required to reduce the institutional characterizations to "nicer," more structural ones; if there were a reduction of institutional types to, say, physiological types, then on Homuncionalism the Identity Theory would be true. Institutional types (at any given hierarchical level of abstraction) are irreducible, though I assume throughout that institutional tokens are reducible in the sense of strict identity, all the way down to the subatomic level.

In fact, the irreducibility of institutional types makes for a mark in favor of Homuncionalism as a philosophical theory of the mental. As Donald Davidson and Wilfrid Sellars have both observed, an adequate theory of mind must, among its other tasks, explain the existence of the mind-body problem itself; this would involve explaining why the mental seems so different from the physical as to occasion Cartesianism in the naive, why it has historically proved so difficult even for the sophisticated to formulate a plausible reduction of the mental to the physical, and why our mental concepts as a family seem to comprise a "seamless whole," conceptually quite unrelated to the physiological or physical family.<sup>9</sup> Homuncionalism provides the rudiments of such explanations. The apparent irreducibility of the mental is the genuine irreducibility of institutional types to the less teleological.<sup>10</sup> The difficulty of outlining a tenable reduction of the mental even to the institutional is due to our ignorance of the organizational

workings of the institution itself at a sufficiently low level of abstraction. Nor is the irreducibility of institutional types to more physiological types an embarrassment, so long as our system of institutional categories, our system of physiological categories, and our system of physical categories are just alternative groupings of the same tokens.

Some philosophers might find the Homuncionalist "reduction" very cold comfort. Certainly it would bore anyone who antecedently understands teleological characterizations of things in terms of mental items such as desires or intentions. Of course, as the foregoing discussion implies, I do not understand teleological talk in that way; rather, I am taking mental types to form a small subclass of teleological types occurring for the most part at a high level of functional abstraction. But if so, then how *do* I understand the teleological?

On this general issue I have little of my own to contribute. I hope, and am inclined to believe, that the teleological characterizations that Homuncionalism requires can be independently explained in evolutionary terms. This hope is considerably encouraged by the work of Karl Popper, William Wimsatt, Larry Wright, Karen Neander, and other philosophers of biology;<sup>11</sup> I cannot improve on their technical discussions. However, I do want to make one theoretical point, and then offer one example to back it up.

The theoretical point is that the teleologicalness of characterizations is a matter of degree: some characterizations of a thing are more teleological than others. One and the same space-time slice may be occupied by a collection of molecules, a piece of very hard stuff, a metal strip with an articulated flange, a mover of tumblers, a key, an unlocker of doors, an allower of entry to hotel rooms, a facilitator of adulterous liaisons, a destroyer of souls. Thus, we cannot split our theory of nature neatly into a well-behaved, purely mechanistic part and dubious, messy vitalistic part better ignored or done away with. And for this reason we cannot maintain that a reduction of the mental to the teleological is no gain in ontological tractability; highly teleological characterizations, unlike naive and explicated mental characterizations, have the virtue of shading off fairly smoothly into (more) brutally physical ones.<sup>12</sup>

Let me give one illustration pertinent to psychology. Consider an organism capable of *recognizing faces* (to take one of Dennett's nice examples of a programmable psychological capacity). There is plenty of point to the question of *how* the organism

does its job; the creature might accomplish its face-recognizing by being built according to any number of entirely dissimilar functional plans. Suppose the particular plan it does use is as follows: It will accept the command to identify only when it is given as input a front view, right profile, or left profile. The executive routine will direct a *viewpoint locator* to look over the perceptual display, and the viewpoint locator will sort the input into one of the three possible orientation categories. The display will then be shown to the appropriate *analyzer*, which will produce as output a coding of the display's content. A *librarian* will check this coded formula against the stock of similarly coded visual reports already stored in the organism's memory; if it finds a match, it will look at the identification tag attached to the matching code formula and show the tag to the organism's *public relations officer*, who will give phonological instructions to the *motor subroutines* that will result in the organism's publicly and loudly pronouncing a name.

Knowing that this is the way in which our particular face-recognizer performs its job, we may want to ask for further details. We may want to know how the viewpoint locator works (is it a simple template?), or how the PR office is organized, or what kinds of subcomponents the analyzer employs. Suppose the analyzer is found to consist of a *projector*, which imposes a grid on the visual display, and a scanner, which runs through the grid a square at a time and produces a binary code number. We may go on to ask how the scanner works, and be told that it consists mainly of a light meter that registers a certain degree of darkness at a square and reports "0" or "1" accordingly; we may ask how the light meter works and be told some things about photosensitive chemicals, etc., etc. Now at what point in this descent through the institutional hierarchy (from *recognizer* to *scanner* to *light meter* to *photosensitive substance*, and as much further down as one might care to go) does our characterization stop being teleological, period, and start being purely mechanical, period? I think it is clear that there is no such point, but rather a finely grained continuum connecting the abstract and highly teleological to the grittily concrete and only barely teleological. And this is why the mental can seem totally distinct and cut off from the physiochemical without being, ontologically, any such thing.<sup>13</sup>

A final word about my reliance on barely explicated teleology: I do not claim that barely explicated

teleology is good or desirable. I do not like it at all, myself. My point is only that the mystery of the mental is *no greater than* the mystery of the heart, the kidney, the carburetor, or the pocket calculator. And as an ontological point it is a very comforting one.<sup>14</sup>

### Advantages of the Teleological Approach

The reader will not have failed to notice that I take *function* very seriously and literally: as honest-to-goodness natural teleology.<sup>15</sup> The policy of taking "function" teleologically has some key virtues: (i) As we have seen, a teleological understanding of "function" helps to account for the perceived *seamlessness* of the mental, the interlocking of mental notions in a way that has nothing visibly to do with chemical and physical concepts.<sup>16</sup> (ii) By imposing a teleological requirement on the notion of functional realization, we avoid all the standard counterexamples to Machine Functionalism, and, I would claim, to any other version of Functionalism; see below. (iii) A teleological functionalism also helps us to understand the nature of biological and psychological *laws*, particularly in the face of Davidsonian skepticism about the latter (Lycan 1981b; Cummins 1983). (iv) If teleological characterizations are themselves explicated in evolutionary terms, then our capacities for mental states themselves become more readily explicable by final cause; it is more obvious why we have pains, beliefs, desires, and so on.<sup>17</sup> (v) The teleological view affords the beginnings of an account of *intentionality* that avoids the standard difficulties for other naturalistic accounts and in particular allows brain states and events to have *false* intentional content. Causal and nomological theories of intentionality tend to falter on this last task (see Lycan 1989).

I have argued above that we need a notion of teleology that comes in degrees, or at least allows for degrees of teleologicalness of characterization, and that we already have such a notion, hard as it may be to explicate – recall the examples of the face-recognizer and the key. Philosophers may differ among themselves as to the correct analysis of this degree notion of teleology – for my own part, I tend to see the degrees as determined by amenability to explanation by final cause, where explanation "by final cause" is reconstructed in turn as a sort of evolutionary explanation (though some

details of this remain to be worked out). But two main points are already clear: (i) At least for single organisms, degrees of teleologicalness of characterization correspond rather nicely to levels of nature.<sup>18</sup> And (ii) there is no single spot *either* on the continuum of teleologicalness or amid the various levels of nature where it is plainly natural to drive a decisive wedge, where descriptions of nature can be split neatly into a well-behaved, purely "structural," purely mechanistic mode and a more abstract and more dubious, intentional, and perhaps vitalistic mode – certainly not any spot that also corresponds to any intuitive distinction between the psychological and the merely chemical, for there is too much and too various biology in between.

My own panpsychist or at least panteleologic tendencies are showing now. Many tougher-minded philosophers will find them fanciful at best, and of course (in my lucid moments) I am prepared to admit that it is hard to see any use in regarding, say, *atomic*-level description as teleological to any degree.<sup>19</sup> certainly explanation-by-final-cause does not persist all the way down. *But*: unmistakably teleological characterization (description that is obviously teleological to some however small degree) persists *as far* down as could possibly be relevant to psychology (well below neuroanatomy, for example). And the *role/occupant* distinction extends much further down still. Thus the vaunted "function"/"structure" distinction as ordinarily conceived by philosophers fails to get a grip on human psychology where it lives. . . .

Everything I have said so far may seem dull and obvious. I hope it does. I am trying to call attention to what I consider a home truth about the structure of the physical world, because I think neglect of this truth, inattention to the hierarchical nature of Nature, has led to significant errors about consciousness and qualia. In what remains I shall briefly discuss a few.

Block (1981), Lycan (1987), and others have put forward various counterexample cases, designed to show that having a functional organization, however complex, is insufficient for hosting qualitative, phenomenally feely states; probably the best known and most discussed of these are Block's "homunculi-head" and "population of China" examples. If such counterexamples are to be rebutted, the Functionalist must exhibit some reasonable requirement that they fail to satisfy,

despite their mimicking in one way or another the functional organization of a real sentient creature.

Homuncualism teleologically understood does the trick with ease. For none of the systems imagined in the counterexamples is teleologically organized in anything like the right way; most are not even organisms at all (see Lycan 1987, chs 3 and 5).

Even if the puzzle cases fail to refute Homuncualism, some problems of chauvinism and liberalism remain to be resolved. Whether or not Fodor and Block are right in suggesting that Putnam moved too far back toward Behaviorism in backing off from the Identity Theory, the Functionalist certainly bears the responsibility of finding a level of characterization of mental states that is neither so abstract or behavioristic as to rule out the possibility of inverted spectrum, etc., nor so specific and structural as to fall into chauvinism. Block himself goes on to argue that this problem is insoluble.

He raises the dilemma for the characterization of *inputs* and *outputs* in particular. Plainly, inputs and outputs cannot be characterized in human neural terms; this would chauvinistically preclude our awarding mental descriptions to machines, Martians, and other creatures who differ from us biologically, no matter what convincing credentials they might offer in defense of their sentience. On the other hand, inputs and outputs cannot be characterized in purely abstract terms (i.e., merely as "inputs" and "outputs"), since this will lead to the sort of ultraliberalism that Block has disparaged by means of his earlier examples and also by means of new ones, such as that of an economic system that has very complex inputs, outputs, and internal states but that certainly has no mental characteristics. Nor can we appeal to any particular sorts of interactions of the sentient being with its environment via inputs and outputs, since in a few cases (those of paralytics, brains *in vitro*, and the like) we want to award mental descriptions to objects that cannot succeed in interacting with their environments in any way. Block concludes,

Is there a description of inputs and outputs specific enough to avoid liberalism, yet general enough to avoid chauvinism? I doubt that there is.

Every proposal for a description of inputs and outputs I have seen or thought of is guilty of either liberalism or chauvinism. Though this

paper has focused on liberalism, chauvinism is the more pervasive problem.

...there will be no physical characterizations that apply to all mental systems' inputs and outputs. Hence, any attempt to formulate a functional description with physical characterizations of inputs and outputs will inevitably either exclude some [possible] systems with mentality or include some systems without mentality.

...On the other hand, as you will recall, characterizing inputs and outputs simply as inputs and outputs is inevitably liberal. I, for one, do not see how functionalism can describe inputs and outputs without falling afoul of either liberalism or chauvinism, or abandoning the original project of characterizing mentality in nonmental terms. I do not claim that this is a conclusive argument against functionalism. Rather, like the functionalist argument against physicalism, it is perhaps best construed as a burden of proof argument.

I am not sure how detailed a plan Block is demanding of the Functionalist here, though I have agreed that, on a mild-mannered understanding of "burden of proof," Block's challenge is one that the functionalist does bear the burden of meeting. The question is whether this burden is as prohibitively heavy as Block seems to assume. And there are at least three factors that I think lighten it considerably and give us some cause for optimism:

First, there is a line of argument that offers at least some slight positive reason or natural motivation for thinking that the dilemma of chauvinism and liberalism (either in regard to inputs and outputs or in regard to the inner states that the Functionalist identifies with our mental states) does admit a solution. It begins as a slippery-slope argument. Block has stated the dilemma very uncompromisingly, implying that one's *only* choices are (a) to characterize inputs and outputs physiologically and be a chauvinist, or (b) to characterize inputs and outputs "purely abstractly" and be a bleeding heart. But this brutal statement of the alternatives overlooks the fact... that functional abstraction is a matter of degree. Purely physiological characterization is an extreme, lying at the lower or "more structural" end of the spectrum; "purely abstract" characterization is the opposite extreme, lying at the higher or "more functional" end. Notice that... there are

characterizations that are even *more* "structural" than physiological ones are, such as microphysical ones, relative to which physiological ones are "functional"; similarly, there are really more abstract characterizations than "input" and "output" themselves, such as "transfer," "motion," or even "occurrence." If it is true, as it seems to be, that "purely abstract" characterizations and physiological characterizations merely lie near the two ends of a continuum of functional abstraction, then it is reasonable to expect that there exists some intermediate level of abstraction that would yield characterizations that rules out the Bolivian economy, the Abnegonian Galaxy, the microbiology of the Everglades and their ilk, but would make room for human beings, molluscs, Martians, and brains *in vitro*. The truth lies (as it so often does) somewhere in between, and, depending on which aspect of which mental state interests one, not always at the same spot in between either. Wait and see what resources will be available at various intermediate levels.<sup>20</sup> . . .

Let us remember in addition (here is my second point in response to Block's challenge) that nothing forces us to assume that all the different kinds of mental states occur at the *same* level of functional abstraction. The intuitively "more behavioral" sorts of mental states, such as beliefs and desires and intentions, presumably occur at a relatively high level of abstraction, and this makes it easy for us to ascribe beliefs and desires and intentions to Martians whose overt behavior and very superficial psychology match ours; the same is true of highly "informational" mental activities such as remembering and (literal) computing. Intuitively, "less behavioral," more qualitative mental states probably occur at a much lower level of abstraction; sensings that have certain particular kinds of qualitative characters probably *are* quite specific to species (at least, we should not be very surprised to find out that this was so), and quite possibly our Martian's humanoid behavior is prompted by his having sensations (or possibly "schmensations") somewhat unlike ours, despite his superficial behavioral similarities to us.

I am not aware that anyone has ever explicitly defended Two-Levelism as such.<sup>21</sup> But Two-Levelism seems to be what lies directly behind such apparent dilemmas as Block's "problem of the inputs and the outputs."

Parallel considerations apply to the problem of intentionality. We think that a state of an organism is either an intentional state or not, period, and

then we wonder what the functional or institutional locus of intentionality might be. I do not think intentionality can be a *purely* functional property at all, for reasons that are now familiar,<sup>22</sup> but insofar as it is, I think we would do well to admit that intentionality itself comes in degrees.<sup>23</sup> The "marks" of intentionality or aboutness are none too clear, but what does seem clear upon reflection is that there is an intermediate level of functional characterization that offers a *kind* of directedness-upon-a-possibly-nonexistent-object-or-type that nevertheless falls short of the rich, full-blooded intentionality exhibited by the human mind. At this intermediate level, we speak systems-theoretically of "detectors," "scanners," "filters," "inhibitors," and the like, meaning these terms quite literally but without actually imputing *thought* or what might be called "occurrent" aboutness. But I must leave the development of these observations for another occasion.<sup>24</sup>

Third, it might be profitable for us simply to stand by the "purely abstract" characterization of inputs and outputs, throwing the whole problem of chauvinism and liberalism back onto our characterization of internal states and events. There are so many possibilities, so many different levels of abstraction in the functional hierarchy as it applies to the brain (many of which overlap and cut across each other), that it seems quite reasonable to expect there to be, for each mental state-type, some middle way between chauvinism and liberalism - not necessarily the *same* middle way for each state-type. It is simply an error to think that all mental phenomena must be functionally located at the same level, or that any single mental state must be localized entirely at one level. Regarding the "more functional," *nearly* behavioristic mental states, perhaps we would not even mind admitting that an economic system or the population of China could have such states (say, dispositional beliefs), if it were to come to that. And possibly at the least functional end of the continuum there are even mental state-types of which the Identity Theory is true, though it is hard to think of any mental state that is as "qualitative" as that.

The foregoing remarks suggest a final additional response to Block's "absent qualia" arguments, one that I think is virtually conclusive. Earlier I characterized Block's intuitive disquiet over Functionalism as being a matter of feeling the incongruity between the relationalness of Functionalist explications and the homogeneous, primitively

*monadic* qualitative characters of their explicanda; I gather that this incongruity seems to him absolute. Notice that evidently he has no similar objection to the Identity Theory; like any other materialist, he would simply charge the Identity Theorist with chauvinism and raise no further complaint. After all, one of the theory's main advantages was its ability to account for the possibility of inverted spectrum or other inner variation despite outward conformity. But if we also accept my claim that Homofunctional characterizations and physiological characterizations of states of persons reflect merely different levels of abstraction within a surrounding functional hierarchy or continuum, then we can no longer distinguish the Functionalist from the Identity Theorist in any absolute way. "Neuron," for example, may be understood either as a physiological term (denoting a kind of human cell) or as a (teleo-) functional term (denoting a relay of electrical charge); on *either* construal it stands for an instantiable - if you like, for a role being played by a group of more fundamental objects. Thus, *even the Identity Theorist is a Functionalist* - one who locates mental entities at a very low level of abstraction. The moral is that if Block does want to insist that Functionalist psychology is stymied by a principled incongruity of the sort I have mentioned and that a philosophy of mind that explicates mental items in terms of relational roles or instantiables cannot in principle accommodate the intractable monadicity of qualia, then one would have to make the same charge against the Identity Theorist as well, and this, I trust, he feels no intuitive compulsion to do.<sup>25</sup> In fact, Block lets that Theory cop a plea of species chauvinism overall, and even allows that it is probably true of some mental properties.

There is an idea, brought on by blind Two-Levelism, that Functionalism differs somehow *conceptually or structurally* from the Identity Theory, in such a way as to incur different sorts of objections. As I have said, the Identity Theory is just an empirically special case of Functionalism, one that (implausibly) locates all mental states at the same very low level of institutional abstraction - the neuroanatomical. Thus there should be no purely conceptual or philosophical objections that apply to Functionalism that do not apply to the Identity Theory or vice versa, even if one is empirically less reasonable than the other. Yet philosophers such as Block have claimed to see such objections. If my doctrine of the continuity of nature is right, something must be wrong here;

for neuroanatomical terms are functional and so relational just as higher-organizational terms are, albeit at a lower level of abstraction. If there is a principled incongruity between relational characterization and the intrinsicness of phenomenal quality, and if that incongruity stymies Functionalism, then it should preclude the Identity Theory as well.<sup>26</sup>

Consider a second example of such an objection: Block further contends that Functionalism is unable to allow the possibility of "inverted spectrum" or other types of internally switched qualia unreflected even counterfactually in behavior - unable in a way that the Identity Theory is not, since the Identity Theory is *made to order* for representing cases of inverted qualia. But if my reflections on the continuity of levels of nature are right, something must be amiss here. And something is. Just as it is easy to imagine undetectably switched *neurophysiology* underlying inverted spectrum (see Lycan 1973), it is easy to imagine a switching of functional components more abstractly described (though doubtless there are limits to this, and quite possibly one could not ascend to a very much higher level of abstraction and keep the inversion behaviorally undetectable).

The truth of the matter is obscured by a pragmatic ambiguity in the notion of "inverted qualia," an ambiguity that I think has lent Block rhetorical aid even though it is far from subtle. To wit, there is a hidden parameter: "inverted" *with respect to what?* (Compare the correlative relation of *supervenience*: supervenient on what?) Traditionally, "inverted spectrum" has meant (color) qualia inverted with respect to actual and counterfactual input/output relations alone. Either from duty or by inclination, Analytical Behaviorists and Wittgensteinians denied the conceivability of *that* inversion, but most people's ordinary modal intuitions have favored it, and Identity and Functionalist theories alike have accommodated it with ease; it has never posed any threat to Functionalism.

What would damage Functionalism is the conceivability of qualia inverted with respect to I-O relations *plus* internal functional organization. This inversion hypothesis is much stronger and more daring. Its possibility is controversial to say the least. Indeed, to assert it is simply to deny the truth of Functionalism - it is to say without argument that two organisms could differ in their qualitative states even though they were exactly alike in their entire global functional organization, at whatever level of institutional abstraction is in

question. Of course there have been philosophers who have insisted without argument on the metaphysical possibility of organisms' differing in their qualitative states despite being *molecular* duplicates, for that matter, but such insistence has no intrinsic credibility even if the relevant theories of mind turn out in the end to be false. The possibility of spectrum inverted with respect to I-O relations alone is a well-entrenched and respectable though I suppose defeasible modal intuition; the possibility of spectrum inverted with respect to I-O relations *plus internal functional organization at however low a level of abstraction proponents feel it plausible to name* is anything but obvious and in conflict with some intuitively plausible supervenience theses.

(Some relationally minded theorists may find it natural to assume a certain *privileged* level of abstraction at the outset. For example, "analytical functionalists," or as I prefer to call them, commonsense relationalists, who hold that the *meanings of mental terms* are determined by the causal roles associated with those terms by common sense or "folk psychology," thereby deny themselves appeal to any level of functional organization lower than is accessible to common sense.<sup>27</sup> Folk psychology aside, the "High Church" computationalists<sup>28</sup> scorn appeal to human biology even within a purely *scientific* account of cognition and behavior, though their own chosen level of nature is none too clearly specified.<sup>29</sup> A theorist who cleaves to such a privileged level of organization may of course admit "inverted spectrum" relative to that chosen level, so long as he or she is willing to type-identify qualia with still lower-level items.<sup>30</sup>

## Two Alternative Strategies

I have recommended one way of solving the problems of chauvinism and liberalism concerning qualia within a Functionalist ontology of the mental. There are alternative possible strategies. One alternative approach would be to bifurcate our view of the mental, by simply taking over the distinction between a mental state and its qualitative character, explicating the states in functional terms and the characters in rather broad physiological terms, tolerating the consequence that inverted spectrum or lesser interpersonal differences in qualia might be more prevalent than we think (*viz.*, exactly as prevalent as are interpersonal

physiological differences of comparable magnitude).<sup>31</sup>

*Pain* would present a useful test case for this second suggested way of accommodating qualia. An interesting and distinctive thing about pain is that (unlike most other mental states) it has both a strongly associated behavior pattern *and* vivid introspectible feely properties. This means, on the present proposal, that pain states may receive *multileveled* analyses. For example (just to speculate a bit), we might end up wanting to classify any internal state of an organism that played pain's usual "gross" behavioral role (that of being caused by damage and producing withdrawal-*cum*-favoring) as being a pain, but to distinguish the feels of pains according to the states' physiological bases.<sup>32</sup> It would follow that, although mollusks and Martians have pains, their pains probably feel differently to them from the ways in which our pains feel to us. It would also follow that a state that feels like a pain state of mine might in a differently organized creature be a mental state of some kind other than pain; some philosophers may find this crassly counterintuitive.

Incidentally, the bifurcated view has become fairly popular in the past few years,<sup>33</sup> and is often expressed by saying that (e.g.) "pain itself is functional while its specific feel is neurophysiological." But the latter formulation again presupposes Two-Levelism. See the "functional"/"structural" distinction as level-relative, and the bifurcated theory collapses into a pointlessly specific version of the thesis (which I hope will become a truism) that mental states and their qualitative characters may well not be explicated in terms of the same level of nature (in particular, the locus of qualitative character may be lower in the hierarchy than that of the mental state generically considered). I emphatically agree with the latter thesis, as I have already indicated, but it is no competing *alternative* to Functionalism.

A third alternative approach suggests itself for the case of bodily sensations (though I doubt whether it could easily be applied to perceptual qualia). It is to suppose that feelings that seem phenomenally to be simple are actually complex and that the distinctive quale associated with a feeling of a certain type is really the coincidence or superimposition of a number of distinct, individually manageable homunctional features. I think this line, rather than that adumbrated in the foregoing paragraph, is the most plausible to take for the case of pain, because it is strongly

suggested by the anesthesiological data collected and summarized by Dennett (1978, ch. 11). What these data seem to indicate is that chemically different anesthetics and analgesics disrupt subjects' normal "pain" subroutines at different functional junctures, eliciting from the subjects quite different verbal reports of their effects. Of a group of subjects suffering pain of roughly the same kind and intensity, one subgroup given drug *A* may report that the pain has diminished or gone away entirely, whereas a subgroup given drug *B* may report that although they know that the pain is still there, they cannot feel it; a subgroup given drug *C* may say that although they can still feel the pain just as intensely as ever, they do not *mind* any more; and so on. That some of these reports sound funny to us (they would be pooh-poohed as "unintelligible" by some Wittgensteinians) naturally reflects the fact that the subjects' normal inner workings are being disrupted, and their normal inner experience of pain being altered, by the drugs. What the drugs seem to be doing is *splitting off components* of the subjects' phenomenal experience of the pain, by splitting off component sub-routines of its rather complicated functional basis. And if this is so, it follows that our phenomenal experience of pain *has* components - it is a complex, consisting (perhaps) of urges, desires, impulses, and beliefs, probably occurring at quite different levels of institutional abstraction. If these components can individually be split off from each other by drugs, then we may perform a *Gedanken-experiment* in which we hypothetically take a suffering subject, split off one component of his pain by administering drug *A*, then split off another component by administering drug *B*, and repeat

this process, eliciting reports as we go to keep track of how we are doing. It seems to me plausible to think that if we were to keep this up, disrupting one access pathway after another and eliminating the component urges, desires, and beliefs one by one, we would sooner or later succeed in eliminating the pain itself; it also seems that if we were to reverse the process - to begin restoring the pathways by withholding the various drugs one by one - the subject would necessarily come to feel the full-fledged pain again (provided his damaged tissues had not been repaired in the meantime). I believe this makes it reasonable to suppose that some (again) *multileveled* proper subsequence of the relevant complex of functional goings-on is both necessary and sufficient for the occurrence of the pain, contrary to the spirit of Block's anti-liberalism.

I do not know how to make a conclusive choice among the three alternative approaches I have described, or what sorts of further evidence we might seek. I have run through some of the options only in order to show that the Homunctionalist has fairly rich resources that can be brought to bear both on the dilemma of chauvinism and liberalism and on the positive task of accounting for qualia. On the basis of these resources I believe we are entitled to conclude that Block's pessimism about qualia is unwarranted. . . .

If my continuity doctrine is obvious as stated, it has not been obvious enough to some of our leading philosophers of mind. I hope the foregoing demonstrations will also serve to make Homunctionalism all the more attractive as a theory of the mental.

## Notes

- 1 I believe just as firmly in some form of act-utilitarianism in ethics, but the sacred principle of utility itself forbids my even telling you this, much less committing (detectable) murders in its name.
- 2 This multileveled hierarchical structure was noted and eloquently presented by Herbert A. Simon (1969); I do not know if the idea predates him. William C. Wimsatt has also written brilliantly on it (1976). Its application to psychology was first brought to my attention by Fodor (1968) and Dennett (this volume); see further references below.
- 3 I have in mind Lewis Thomas's (1974) discussion of insect societies and of the relation between

human beings and their own mitochondria. The micrology of "organisms" is highly interest-relative. Note well, we must grant a pluralism of different reductive relations between levels of nature; consider also the entirely tenable notion of the *corporation as person* (Biro 1981; French 1984; Brooks 1986).

- 4 For a rich exposition and defense of the strategy, see Cummins (1983). However, Richardson (1983) throws some fairly cold water.
- 5 For philosophically relevant discussion and references, see P. M. Churchland (1986) and P. S. Churchland (1986).
- 6 Dennett's main concern in the work containing the following passage is the explication of intentionality.

- That concern is not mine here; I am interested only in homuncular breakdown *per se*.
- 7 In fact, as David Armstrong has pointed out to me, the present maneuver blocks a number of typical infinite-regress arguments in the philosophy of mind, including Ryle's complaint against volitional theories of deciding. Dennett himself wields it against "Hume's problem" regarding self-understanding representations (1978, pp. 122ff.).
  - 8 For an actual hands-on homuncular breakdown of the speech center, see figure 1, p. 262 of Lycan (1984).
  - 9 For stout insistence of this, see Davidson (1970).
  - 10 Thus, Smart's example of the logic of "nation" statements' being different from the logic of "citizen" statements may have been more apropos than he imagined.
  - 11 Popper (1972); Wimsatt (1972); Wright (1973); Millikan (1984); Neander (1981, 1983). Neander's evolutionary explication is the best I know. It is criticized with effect by E. Prior in an unpublished note and by Pargetter and Bigelow (1987); the truth seems to me to lie somewhere in between. Jonathan Bennett (1976) offers a different naturalistic approach to teleology due to Ann Wilbur MacKenzie (1972) (and in discussion has urged me to switch).
  - 12 Characterizations of the contents of our space-time slice may thus be arranged in a continuum, from the least teleological to the most (highly) teleological. This continuum corresponds fairly neatly to the hierarchy of functional instantiation or realization. The molecules jointly realize, or play the role of, the piece of metal; the piece of metal plays the role of the key; the key serves as our door-unlocker; and so on. The prevalence of functional hierarchies of this kind, I believe, is what encourages ontological reduction and the idea that "everything is ultimately a matter of physics." On the relations between teleology viewed from an evolutionary perspective, functional hierarchies, ontology, and the methodology of scientific reduction, see again Wimsatt (1976). I have also profited from reading Mellick (1973), and see Matthen and Levy (1984).
  - 13 As Jerry Fodor has pointed out to me in discussion, there is one tolerably clear distinction that a Two-Leveler might have in mind and that is absolute: it is the distinction between objects whose proper parts are essential to them and objects whose parts are not. For example, a bicycle's or even a tree's parts are replaceable, while a water molecule's parts perhaps are not (one might argue either that if the molecule were to lose one of its hydrogen or oxygen atoms it would not be *that* molecule or that without the right sorts of atoms it would not be a water molecule at all). I agree that this distinction is gen-

uine, and I expect it has some metaphysical importance. But it has no *psychological* importance. The level of chemistry is far too low in the institutional hierarchy to affect mentation; that is, if two neuro-anatomies are just the same even though they are realized by different chemicals, psychology is the same.

- 14 Amelie Rorty has suggested to me the Aristotelian idea of explaining an organism's component functions (more exactly, of explaining its-functions'-constituting-its-thriving) by reference to the suitability of those functions for the material conditions of the organism's species. This idea fits well with the etiological account of function that I tend to favor. Given a relatively undifferentiated mass of "lower" biological material at a much earlier evolutionary stage, how would it clump together and articulate itself in order to face the world at large in a more robust and less vulnerable way? Its own "structural" or "material" nature would enforce some answers and suggest still others, and given selection pressures of various now retrodictable sorts it is no surprise that many or most of these answers have been realized. If "function" is understood in evolutionary terms, then, function itself gets explained in this way, in terms of the propensities of the organism's material substratum. I take that explanation to complement, rather than to compete with, "downward-causation" explanations based in higher levels of nature (of the sort Wimsatt talks about). In fact, we get a sort of pincer movement: selection pressure from much higher levels interacting with bottom-up pressure from the nature and propensities of the particular chemical constitution of the pre-existing neighborhood, the two pressures jointly molding what lies between. But one might want to emphasize the bottom-up pressure at the expense of higher-level explanations. *In some sense* that emphasis has to be right, given supervenience of top on bottom, though it is tricky to work out all the different up-down interrelations there are.

Rorty points out (in correspondence) that full-scale multiple realizability must be distinguished from mere functional characterization of states of organisms, since detailed accounts of function tend to put strict requirements on realizing-stuff; there is a trade-off here. But I do not see that the Aristotelian bottom-up explanation strategy *per se* counts against multiple realizability. For the same functional answers or solutions might well be hit upon by chemically quite different bunches of primordial stuff. Rorty offers the example of *eating*: Computers do not eat, in any literal sense, and the earth does not ingest rain; multiple realizability fails even though the activity is functionally characterized. I want to make the same sort of rejoinder that I shall be making to an argument of Block's below: Of

course computers and other (even biologic) entities do not eat; but there is an intermediate, more abstract characterization of eating itself *holotropism* as it was called in my college biology classes — which excludes computers but includes lots of species biochemically quite different from ours; it has something to do with acquiring proteins very similar to one's own and physically homogenizing them and ingesting them and making them part of one without major rearrangement of amino acids or something of the sort — at any rate, it is a form of nourishment that is sharply distinguished from many other species' and is rather distinctive of our phylum or whatever. This point checks nicely with my usual idea of functional characterizations that hold for intermediate levels of nature and are neither too vague and general nor too chauvinistically species-bound.

- 15 Elliott Sober (this volume) praises this attitude as "putting the function back into functionalism"; cf. my remarks on p. 27 of Lycan (1981a) regarding Putnam and Fodor's pun on the word "function."
- 16 For details, see Lycan (1981b).
- 17 Why does pain hurt? Why could we not have a damage-signaling and repair-instigating system that was not uncomfortable? The answer is simple. Suppose I had just such a system, like the red warning light on my auto engine. Just as I habitually though irrationally ignore the warning light and vaguely hope it will go away, I would ignore a personal warning light if it did not intrinsically provide me with an urgent motive to do something about it.
- 18 Robert Van Gulick has presented me (in correspondence) with some meteorological and geological cases in which (apparent) degrees of teleologicalness do not follow levels of nature. Such cases are very much to the point, but I shall have to postpone going into them.
- 19 Ned Block, who violently disagrees with me on the present issues, once said (in conversation), "I'll give you *neurons* and *cells* and so on as functional, but when you come to *hydrogen* and *oxygen*, when you get right down to the level of *chemistry*, there's just nothing functional or teleological at all!" Oh, no? "Hydro-"*what?* "Oxy-"*what?* (The shot is a cheap one but immensely satisfying.)
- 20 "Wait till next year!" John Searle jeers in a different but very similar connection (1980). *Of course* wait till next year!
- 21 Save perhaps the "analytical functionalist," whose view I reject (see note 27).
- 22 Putnam (1975); Fodor (1980); Stich, this volume; Burge (1979); Lycan (1981c); ...
- 23 This idea is anticipated in part by Dretske (1981). See also Van Gulick (this volume, 1982).
- 24 I would also observe that some current disputes within the cognitive science community are miscon-

ceived in the Two-Levelist way. For example, the "bottom-uppers" versus the "High-Church Computationalists" (see P. S. Churchland 1986, and Dennett 1986) and the New Connectionists versus the same (see Bechtel 1985). The New Connectionists in particular are a superb example of a biocomputational middle way. Somewhat in the same spirit is P. M. Churchland's (1986) "phase-space sandwich" model of sensorimotor coordination, based on Pelhousz and Linas (1979, 1982); or rather, though he does not always think it in a mediating way, I count it as another feasible middle way *within the spirit of a properly teleologized functionalism*.

- 25 Wilfrid Sellars does. But that is another story ...; see chapter 8 of Lycan (1987).
- 26 Block does not himself stress the relational/monadic contrast, but offers his differential intuitions raw; so he may remain unmoved by my foregoing *ad hominem* and simply insist that having a *neurochemistry* roughly like ours is a necessary condition for experiencing qualia, relational or not. Yet, I wonder, how could a *philosopher* know *that?* Is it aglow with the Natural Light?
- 27 I am indebted to Sydney Shoemaker for useful correspondence on this point. For my own part, I cannot accept analytical functionalism, for two reasons: (i) I reject the alternatively conceptual-analysis or implicit-definition theory of meaning on which that theory rests. (See Armstrong 1968 and Lewis 1972 for its two most explicit versions and defenses, and Lycan 1981b, especially note 10, for my alternative view of the semantics of mental terms; also, for a similar view, see Jacoby 1985.) (ii) I doubt that common sense or "folk psychology" contains enough information about mental entities to characterize their natures as richly as would be needed to avoid counterexample. Clothespin models of folk psychology would be pretty easy to come by, without the massive complexity and teleological organization that would warrant an ascription of real mentality.
- 28 The term is due to Dennett (1986).
- 29 Here I follow some recent writers in supposing that there are really any High Church Computationalists; I am not sure that any actual Functionalists has ever self-consciously intended the view. It is usually ascribed to Zenon Pylyshyn and Jerry Fodor, on the basis of some of their remarks about multiple realizability. Perhaps Ned Block does really hold it, or he would not continue to resist my case against Two-Levelism as begun in my 1981a.
- 30 See particularly (again) Bechtel (1985) and the references made therein.
- 31 Block hints on p. 460 that he might not find this suggestion entirely congenial. And see note 33.
- 32 This move would take some of the sting out of what I take to be an anti-Functionalist argument in David Lewis (1980).

- 33 Block hinted at this view, as I have mentioned. I developed the suggestion in Lycan (1981a), pp. 47-8. It has also been picked up by Hilary Putnam (1981), Sydney Shoemaker (1981), Patricia Kitcher

## References

- Armstrong, D. M. (1968) *A Materialist Theory of the Mind*, Routledge & Kegan Paul.
- Attneave, F. (1960) "In defense of homunculi," in W. Rosenblith (ed.), *Sensory Communication*, MIT Press.
- Bechtel, P. W. (1985) "Contemporary connectionism: Are the new parallel distributed processing models cognitive or associationist?" *Behaviorism* 13, 53-61.
- Bennett, J. (1976) *Linguistic Behaviour*, Cambridge University Press.
- Bigelow, J. and Pargetter, R. (1987) "Functions," *Journal of Philosophy* 84, 181-96.
- Biro, J. (1981) "Persons as corporate entities and corporations as persons," *Nature and System* 3, 173-80.
- Block, N. J. (1981) "Psychologism and behaviorism," *Philosophical Review* 90, 5-43.
- Brooks, D. H. M. (1986) "Group minds," *Australasian Journal of Philosophy* 64, 456-70.
- Burge, T. (1979) "Individualism and the mental," in P. French, T. E. Uehling and H. Wettstein (eds), *Midwest Studies in Philosophy, vol. IV: Studies in Metaphysics*, University of Minnesota Press.
- Churchland, P. M. (1986) "Some reductive strategies in cognitive neurobiology," *Mind* 95, 223-38.
- Churchland, P. S. (1986) *Neurophilosophy*, Bradford Books/MIT Press.
- Cummins, R. (1983) *The Nature of Psychological Explanation*, Bradford Books/MIT Press.
- Davidson, D. (1970) "Mental events," in L. Foster and J. W. Swanson (eds), *Experience and Theory*, University of Massachusetts Press.
- Dennett, D. C. (1978) *Brainstorms*, Bradford Books.
- Dennett, D. C. (1986) "The logical geography of computational approaches: a view from the East Pole," in M. Brand and R. M. Harnish (eds), *The Representation of Knowledge and Belief*, University of Arizona Press.
- Dretske, F. (1981) *Knowledge and the Flow of Information*, Bradford Books/MIT Press.
- Fodor, J. A. (1980) "Methodological solipsism considered as a research strategy in cognitive psychology," *Behavioral and Brain Sciences* 3, 63-73.
- French, P. (1984) *Collective and Corporate Responsibility*, Columbia University Press.
- Horgan, T. (1984) "Functionalism, qualia, and the inverted spectrum," *Philosophy and Phenomenological Research* 44, 453-70.
- (1982), Terence Horgan (1984), and Gregory Sheridan (1986) among others (Shoemaker calls its "selective parochialism").
- Jacoby, H. (1985) "Eliminativism, meaning, and qualitative states," *Philosophical Studies* 47, 257-70.
- Kitcher, P. (1982) "Two versions of the Identity Theory," *Erkenntnis* 17, 213-28.
- Lewis, D. (1972) "Psychophysical and theoretical identifications," *Australasian Journal of Philosophy* 50, 249-58.
- Lewis, D. (1980) "Mad pain and Martian pain," in N. Block (ed.), *Readings in Philosophy of Psychology*, vol. 1, Harvard University Press.
- Lycan, W. (1973) "Inverted spectrum," *Ratio* 15, 315-19.
- Lycan, W. (1981a) "Form, function, and feel," *Journal of Philosophy* 78, 24-50.
- Lycan, W. (1981b) "Psychological laws," *Philosophical Topics* 12, 9-38.
- Lycan, W. (1981c) "Toward a homuncular theory of believing," *Cognition and Brain Theory* 4, 139-59.
- Lycan, W. (1984) *Logical Form in Natural Language*, Bradford Books/MIT Press.
- Lycan, W. (1987) *Consciousness*, Bradford Books/MIT Press.
- Lycan, W. (1989) "Ideas of representation," in *Mind, Value, and Culture: Essays in Honor of E. M. Adams* (edited by D. Weissbord), Ridgeview Publishing.
- MacKenzie, A. W. (1972) "An analysis of purposive behavior," Cornell University doctoral dissertation.
- Matthen, M. and Levy, E. (1984) "Teleology, error, and the human immune system," *Journal of Philosophy* 81, 351-71.
- Mellick, D. (1973) "Behavioral strata," Ohio State University doctoral dissertation.
- Millikan, R. G. (1984) *Language, Thought, and Other Biological Categories*, Bradford Books/MIT Press.
- Neander, K. (1981) "Teleology in biology," Unpublished Xerox.
- Neander, K. (1983) "Abnormal psychobiology," La Trobe University doctoral dissertation.
- Pellionisz, A. and IJinas, R. (1979) "Brain modelling by tensor network theory and computer simulation. The cerebellum: distributed processor for predictive coordination," *Neuroscience* 4, 323-48.
- Pellionisz, A. and IJinas, R. (1982) "Space-time representation in the brain. The cerebellum as a predictive space-time metric tensor," *Neuroscience* 7, 2,949-70.
- Popper, K. (1972) "Of clouds and clocks: an approach to the problem of rationality and the freedom of man," in *Objective Knowledge: An Evolutionary Approach*, Oxford University Press.
- Putnam, H. (1975) "The meaning of 'meaning'," in K. Gunderson (ed.), *Minnesota Studies in the Philosophy of Science, vol. 7: Language, Mind and Knowledge*, University of Minnesota Press.
- Putnam, H. (1981) *Reason, Truth and History*, Cambridge University Press.
- Richardson, R. (1983) "Computational models of mind," Unpublished monograph.
- Searle, J. (1980) "Minds, brains and programs," *Behavioral and Brain Sciences* 3, 417-24.
- Sheridan, G. (1983) "Can there be moral subjects in a physicalistic universe?" *Philosophy and Phenomenological Research* 43, 425-48.
- Sheridan, G. (1986) "Selective parochialism and Shoemaker's argument for functionalism," typescript, Western Michigan University.
- Shoemaker, S. (1981) "Some varieties of functionalism," *Philosophical Topics* 12, 93-119.
- Simon, H. (1969) "The architecture of complexity," in *The Sciences of the Artificial*, MIT Press.
- Thomas, L. (1974) *Lines of a Cell*, Bantam Books.
- Van Gulck, R. (1982) "Mental representation - a functionalist view," *Pacific Philosophical Quarterly* 63, 3-20.
- Wimsatt, W. C. (1972) "Teleology and the logical structure of function statements," *Studies in History and Philosophy of Science* 3, 1-80.
- Wimsatt, W. C. (1976) "Reductionism, levels of organization, and the mind-body problem," in G. Globus, G. Maxwell, and I. Savodnik (eds), *Consciousness and the Brain*, Plenum.
- Wright, L. (1973) "Functions," *Philosophical Review* 82, 139-68.