

## AGAINST MATERIALISM

Alvin Plantinga

Materialism is the idea that human beings are material objects—brains, perhaps, or some part of the brain—without immaterial selves or souls. I give two arguments against materialism. The first is an argument from possibility: it is possible that I should exist when no part of my body exists, in which case I am not identical with my body or any part of my body. The second is an argument from impossibility. If materialism is true, a belief, for example, the belief that the British lost the American Revolutionary War, would have to be something like an event in one's nervous system. But it is impossible that a material process or event be about something; no such process could be about the British or anything else. So either there are no beliefs or materialism is false. But obviously there are beliefs.

I propose to give two arguments against materialism—or, if you think that's too negative, two arguments for substantial dualism. 'Substantial' is to be taken in two senses: first, the dualism in question, the dualism for which I mean to argue, is substantial as opposed to trivial; some versions of property dualism seem to me to be at best wholly insubstantial. Second, according to the most popular form of dualism—one embraced by Plato, Augustine, Descartes and a thousand others—a human person is an immaterial substance: a thing, an object, a substance, a suppositum (as my Thomist colleagues would put it), and a thing that isn't material, although, of course, it is intimately connected with a material body. But there is also the view the name 'dualism' suggests: the view according to which a human person is somehow a sort of composite substance S composed of a material substance S\* and an immaterial substance S\*\*.<sup>1</sup> We can sensibly include this view under 'dualism'—provided, that is, that having S\* as a part is not essential to S. (I add this proviso because my first argument is for the conclusion that possibly, I exist when my body does not.)

Perhaps a better name for the view I mean to defend is ‘immaterialism’; the view that a human person is not a material object. Of course it’s far from easy to say just what a material object is.<sup>2</sup> For present purposes let’s put it recursively: a material object is either an atom, or is composed of atoms. Thus atoms, molecules, cells, hearts, brains and human bodies are all material objects; we’ll leave open the question whether such things as electrons, quarks, protons, fields, and superstrings (if indeed there are such things) are material objects. What I’ll argue for, accordingly, is the view that human persons are not material objects. They are objects (substances), however; therefore they are immaterial objects. My conclusion, of course, is hardly original (going back at least to Plato); my general style of argument also lacks originality (going back at least to Descartes and possibly Augustine). But the method of true philosophy, unlike that of liberal theology and contemporary French thought, aims less at novelty than at truth.

Three more initial comments: (i) when I speak of possibility and necessity, I mean possibility and necessity in the broadly logical sense—metaphysical possibility and necessity, as it is also called. (ii) I won’t be arguing that it is possible that I (or others) can exist disembodied, with no body at all.<sup>3</sup> (iii) I will make no claims about what is or isn’t conceivable or imaginable. That is because imaginability isn’t strictly relevant to possibility at all; conceivability, on the other hand, is relevant only if ‘it’s conceivable that p’ is to be understood as implying or offering evidence for ‘it’s possible that p’. (Similarly for ‘it’s inconceivable that p’.) It is therefore simpler and much less conducive to confusion to speak just of possibility. I take it we human beings have the following epistemic capacity: we can consider or envisage a proposition or state of affairs and, at least sometimes, determine its modal status—whether it is necessary, contingent, or impossible—just by thinking, just by an exercise of thought.<sup>4</sup>

### I. The Replacement Argument: An Argument from Possibility

I begin by assuming that there really is such a thing, substance or suppositum as I, I myself. Of course I'm not unique in that respect; you too are such that there really is such a thing as you, and the same goes for everybody else. We are substances. Now suppose I were a material substance: which material substance would I be? The answer, I should think, is that I would be my body, or some part of my body, such as my brain or part of my brain. Or perhaps I would be something more exotic: an object distinct from my body that is constituted from the same matter as my body and is colocated with it.<sup>5</sup> What I propose to argue is that I am none of those things: I am not my body, or some part of it such as my brain or a hemisphere or other part of the latter, or an object composed of the same matter as my body (or some part of it) and colocated with it. (I'll call these 'eligible' material objects.) For simplicity (and nothing I say will depend on this simplification) I shall talk for the most part just about my body, which I'll name 'B'. (I was thinking of naming it 'Hercules' or maybe 'Arnold', but people insisted that would be unduly self-congratulatory.)

The general strategy of this first argument is as follows. It seems possible that I continue to exist when B, my body, does not. I therefore have the property possibly exists when B does not. B, however, clearly lacks that property. By Leibniz' Law, therefore (more specifically, the Diversity of Discernibles), I am not identical with B. But why think it possible that I exist when my body does not? Strictly speaking, the replacement argument is an argument for this premise. Again, I conduct the argument in the first person, but naturally enough the same goes for you (although of course you will have to speak for yourself).

So first, at a macroscopic level. A familiar fact of modern medicine is the possibility and actuality of limb and organ transplants and prostheses. You can get a new heart, liver, lungs;

you can also get knee, hip, and ankle replacements; you can get prostheses for hands and feet, arms and legs, and so on. Now it seems possible—possible in that broadly logical sense—that medical science should advance to the point where I remain fully dressed and in my right mind (perhaps reading the South Bend Tribune) throughout a process during which each of the macroscopic parts of my body is replaced by other such parts, the original parts being vaporized in a nuclear explosion—or better, annihilated by God. But if this process occurs rapidly—during a period of 1 microsecond, let's say—B will no longer exist. I, however, will continue to exist, having been reading the comic page during the entire process.

But what about my brain, you ask—is it possible that my brain be replaced by another, the brain I now have being destroyed, and I continue to exist? It certainly seems so. Think of it like this. It seems possible (in the broadly logical sense) that one hemisphere of my brain be dormant at any given time, the other hemisphere doing all that a brain ordinarily does. At midnight, we can suppose, all the relevant 'data' and 'information' is 'transferred' via the corpus callosum from one hemisphere—call it 'H<sub>1</sub>'—to the other hemisphere—H<sub>2</sub>—whereupon H<sub>2</sub> takes over operation of the body and H<sub>1</sub> goes dormant. This seems possible; if it were actual, it would also be possible that the original dormant half, H<sub>2</sub>, be replaced by a different dormant half (in the same computational or functional state, if you like) just before that midnight transfer; then the transfer occurs, control switches to the new H<sub>2</sub>, and H<sub>1</sub> goes dormant—at which time it is replaced by another hemisphere in the same computational or functional condition. In a period of time as brief as you like, therefore, both hemispheres will have been replaced by others, the original hemispheres and all of their parts annihilated by God. Throughout the whole process I serenely continue to read the comics.

This suffices, I think, to show that it's possible that I exist when neither my body nor any part of it exists. What about material objects distinct from my body and its parts, but colocated with it (or one of them) and constituted by the same matter as they? I doubt very much that there could be any such things. If objects of this kind are possible, however, the above argument also shows or at least suggests that possibly, I exist when none of them does. For example, if there is such a thing as the matter of which B is composed—if that phrase denotes a thing or object<sup>6</sup>—it too would be destroyed by God's annihilating all the parts of my body.

Of course very many different sorts of object of this kind—object constituted by the matter of my body and colocated with it—have been suggested, and I don't have the space here to deal with them all. However, we can offer a version of the replacement argument that will be relevant to many of them. Turn from macroscopic replacement to microscopic replacement. This could go on at several levels: the levels of atoms, molecules, or cells, for example. Let's think about it at the cellular level. It seems entirely possible that the cells of which my body is composed be rapidly—within a microsecond or two—replaced by other cells of the same kind, the original cells being instantly destroyed. It also seems entirely possible that this process of replacement take place while I remain conscious, thinking about dualism and marveling at some of the appalling arguments against it produced by certain materialists.<sup>7</sup> Then I would exist at a time at which B did not exist.

But is it really true that this process of replacement would result in the destruction of B? After all, according to current science, all the matter in our bodies is replaced over a period of years, without any obvious compromise of bodily integrity or identity. As a matter of fact, so they say, the matter in our brains is completely replaced in a much shorter time.<sup>8</sup> Why should merely accelerating this process make a difference?<sup>9</sup>

Well, as they say, speed kills. When a cell is removed from an organism and replaced by another cell, the new cell doesn't become part of the organism instantaneously; it must be integrated into the organism and assimilated by it.<sup>10</sup> What does this assimilation consist in? A cell in a (properly functioning) body is involved in a network of causal relations; a neuron, for example, emits and responds to electrical signals. A cell receives nourishment from the blood, and cooperates with other cells in various causal activities. All these things take time—maybe not much time, but still a certain period of time. At the instant the new part<sup>11</sup> is inserted into the organism, and until it has begun to play this causal role (both as cause and effect), the new part is not yet a part of the organism, but a foreign body occupying space within the spatial boundaries of the organism. (Clearly not everything, nor even everything organic, within the spatial boundaries of your body is part of your body: think of the goldfish you just swallowed, or a tapeworm.) Let's use the phrase 'assimilation time' to denote the time required for the cell to start playing this causal role. The assimilation time is the time required for the cell to become assimilated into the body; before that time has elapsed the cell is not yet part of the body. To be rigorous, we should index this to the part (or kind of part) and the organism in question; different parts may require different periods of time for their assimilation by different organisms. For simplicity, though, let's assume all parts and organisms have the same assimilation time; this simplification won't make any difference to the argument.

That a given part and organism are such that the time of assimilation for the former with respect to the latter is  $\underline{dt}$ , for some specific period of time  $\underline{dt}$ , is, I take it, a contingent fact. One thinks the velocity of light imposes a lower limit here, but the time of assimilation could be much greater. (For example, it could depend on the rate of blood flow, the rate of intracellular

transport, and the rate at which information is transmitted through neuron or nerve.) God could presumably slow down this process, or speed it up.

There is also what we might call ‘the replacement time’: the period of time from the beginning of the replacement of the first part by a new part to the end of the time of the replacement of the last part (the last to be replaced) by a different part. The time of replacement is also, of course, contingent; a replacement can occur rapidly or slowly. Presumably there is no non-zero lower limit here; no matter how rapidly the parts are replaced, it is possible in the broadly logical sense that they be replaced still more rapidly.

What’s required by the Replacement Argument, therefore, (or at any rate what’s sufficient for it) is

(Replacement) It is possible that: the cells in B are replaced by other cells and the originals instantly annihilated while I continue to exist; and the replacement time for B and those cells is shorter than the assimilation time.

### **Objections and Replies**

(1) Doesn’t a Star Trek scenario seem possible, one in which you are beamed up from the surface of a planet to an orbiting spacecraft, both you, and in this context more importantly, your body surviving the process? This objection is relevant to the Replacement Argument, however, only if in this scenario your body survives a process in which its matter is replaced by other matter, the original matter being annihilated. But that’s not how the Star Trek scenario works: what happens instead is that the matter of which your body is composed is beamed up (perhaps after having been converted to energy), not annihilated. You might think of this case as one of disassembly (and perhaps conversion into energy) and then reassembly. Perhaps your

body could survive this sort of treatment; what I claim it can't survive is the rapid replacement of the matter in question by other matter, the original matter being annihilated.

(2) I've been assuming that you and I are objects, substances; but that assumption may not be as innocent as it looks. Might I not be an event<sup>12</sup>—perhaps an event like a computer's running a certain program? We ordinarily think of an event as one or more objects  $O_1, \dots, O_n$ , exemplifying a property P or relation R, (where P or R may be complex in various ways and, may of course entail extension over time). Perhaps what I am is an event involving (consisting in) many material objects (organs, limbs, cells, etc.) standing in a complex relation. Then, although I wouldn't be a material object, I would be an event involving nothing but material objects—a material event, as we might call it; and why wouldn't that be enough to satisfy the materialist?

Further: suppose I were a material event: why couldn't that event persist through arbitrarily rapid replacement of the objects involved in it? Think of an event such as a battle; clearly there could be a battle in which the combatants were removed and replaced by other combatants with extremely great rapidity. Let's suppose the commanding officer has an unlimited number of troops at his command. He needs 1000 combatants at any given time: eager to spread the risk, he decrees that each combatant will fight for just 30 seconds and then be instantly replaced by another combatant. (Imagine that technology has advanced to the point where the obvious technical problems can be dealt with.) The battle, we may suppose, begins on Monday morning and ends Tuesday night; this one event, although no doubt including many subevents, lasts from Monday morning to Tuesday night—and this despite the constant and rapid replacement of the combatants. Although there are never more than 1000 troops in the field at any one time, several million are involved in the event, by virtue of those rapid replacements. Of



course the replacement could be much faster; indeed, there is no logical limit on the rapidity of replacement of the combatants, the same event (i.e., the battle) persisting throughout. More generally:

- (a) For any duration  $d$  and event  $E$  and substances  $S_1, S_2, \dots, S_n$  involved in  $E$ , if  $S_1, S_2, \dots, S_n$  are replaced by substances  $S_{n+1}, S_{n+2}, \dots, S_{n+n}$  during  $d$ , then there is an event  $E^*$  that persists through  $d$  and is such that at the beginning of  $d$ ,  $E^*$  involves  $S_1, S_2, \dots, S_n$ , and at the end of  $d$  does not involve  $S_1, S_2, \dots, S_n$ , but does involve  $S_{n+1}, S_{n+2}, \dots, S_{n+n}$ .

So events have a certain modal flexibility along this dimension.<sup>13</sup> Now suppose I were an event. Why couldn't the event which I am persist through arbitrarily rapid replacement of the material objects involved in it? Is there any reason, intuitive or otherwise, to suppose not? Perhaps a material substance can't survive the arbitrarily rapid replacement of its parts; is there any reason to think a material event suffers from the same limitation?

(3) We can conveniently deal with objection (2) by considering it together with another. According to Peter van Inwagen, human beings are material objects; a material object, furthermore is either an elementary particle or a living being. Living beings comprise the usual suspects: organisms such as horses, flies, and oak trees, but also cells (neurons, for example), which may not rise to the lofty heights of being organisms, but are nonetheless living beings. It is living horses, flies etc., that are objects or substances. Indeed, 'living horse' is a pleonasm. On van Inwagen's view, there aren't any dead horses; a 'dead horse', strictly speaking, is not really a thing at all and a fortiori not a horse; it is instead a mere heap or pile of organic matter. Once that horse has died, its remains (as we say in the case of human beings) are a mere assemblage of elementary particles related in a certain way; there is no entity or being there in

addition to the particles. A living horse, on the other hand, is a thing, a substance, in its own right and has as parts only other living beings (cells, e.g.) and elementary particles. Strictly speaking, therefore, there isn't any such thing as a hand, or arm or leg or head; rather, in the place we think of as where the hand is, there are elementary particles and other living things (cells, e.g.) related in a certain way.

But by virtue of what is this horse a thing or a substance: under what conditions does an assemblage of elementary particles constitute a thing, i.e., become parts of a substance? When those particles are involved in a certain complex event: a life. Elementary particles can stand in many relations and be involved in many kinds of events; among these many kinds of events are lives; and when elementary particles are involved in that sort of event, then they become parts of a substance. Further, the object, that living thing, exists when and only when the event which is its life exists or occurs. Still further (and here we may be taking leave of van Inwagen) the survival and identity conditions of the organism are determined by the survival and identity conditions of that event, that life. Consider an organism  $O$  and its life  $L(O)$ . The idea is that  $O$  exists in just those possible worlds in which  $L(O)$  occurs; more precisely,  $O$  and  $L(O)$  are such that for any world  $W$  and time  $t$ ,  $O$  exists in  $W$  at  $t$  if and only if  $L(O)$  exists at  $t$  in  $W$ . Hence

(b) Given an organism  $O$  and the event  $L(O)$  that constitutes its life, necessarily,  $O$  exists at a time  $t$  just if  $L(O)$  occurs at  $t$ .

(We can think of 'exists' as short for 'exists, did exist, or will exist'; similarly for 'occurs'.)

This elegant position certainly has its attractions. It's not wholly clear, of course, that there are any elementary particles (perhaps all particles are composed of other particles so that it's composition all the way down, or perhaps what there really is, is 'atomless gunk' configured in various ways<sup>14</sup>); perhaps electrons, etc., aren't particles at all, but perturbances of fields; and

it's a bit harsh to be told that there really aren't any such things as tables and chairs, automobiles and television sets. Nevertheless van Inwagen's view is attractive. Now suppose we add (b) to van Inwagen's view; the resulting position suggests an objection to the Replacement Argument (an objection that doesn't have van Inwagen's blessing). For (again) why couldn't the event which is my life persist through arbitrarily rapid replacement of the objects it involves? Is there any intuitive support for the thought that there is a lower limit on the rapidity of replacement through which this event could persist? If not, then even if I couldn't be a material substance, I could be a material event; no doubt the materialist would find this materialism enough.

We can respond to these two objections together. According to objection (2), I can sensibly think of myself as an event: presumably the event that constitutes my life. Now perhaps the objector's (a) is true: for any replacement, no matter how rapid, there will be an event of the sort (a) suggests. But of course nothing follows about the modal properties of any particular event. So suppose I am an event: nothing about my modal properties follows from or is even suggested by (a); and it is my modal properties that are at issue here. In particular, it doesn't follow that if I were my life, then I could have continued to exist (or occur) through the sort of rapid replacement envisaged in the Replacement Argument. Now turn to (3). Suppose for the moment we concede (b): we still have no reason to think my life, that particular event, the event which is in fact my life, could have survived those rapid replacements of the objects involved in it. No doubt for any such replacement event, there is an event of the sort suggested by (a); nothing follows with respect to the modal properties of the event which is in my life. In particular it doesn't follow that it could have persisted through the sort of rapid replacements we've been thinking about.

So (a) is really a red herring. But there is a more decisive response here. Objection (3) endorses (b), the claim that there is an event—my life—such that, necessarily, I exist just when it does. Objection (2) also (and trivially) entails (b); if I just am my life then, naturally enough, (b) is true. Fortunately, however, (b) is false. For (b) entails

(c) I and my life are such that necessarily, I exist just when it occurs,  
and (as I'll now argue) (c) is false.

Why think (c) is false? First, it's far from clear just which properties events have essentially. Some think it essential to any event that it include just those objects that it does in fact include, and also that these objects exemplify just the properties and relations they do in fact exemplify. If that were true, an event involving an object O's having a certain property could not have occurred if O had not had that property. But that seems a bit strong; surely the Civil War, for example, (that very event) could have taken place even if a particular Confederate soldier had not trodden on a blade of grass he did in fact step on. Still, there are serious limits here. Perhaps the Civil War (the event which is the Civil War) would have existed even if that soldier hadn't trampled that blade of grass; but the Civil War (that event) could not have lasted only 10 minutes. There is a possible world in which there is a very short war between the states (and it could even be called 'The Civil War'); but there is no possible world in which the war that did in fact take place occurs, and lasts for only 10 minutes. Similarly for my life (call it 'L'): if (b) is true, then of course L has existed exactly as long as I have. L, therefore, has by now existed for more than 70 years. Clearly enough, however, I could have existed for a much shorter time: for example, I could have been run over by a Mack truck at the age of six months (and not been subsequently sustained in existence by God). L, however, could not have existed or occurred for only those first few months, just as the Civil War could not have existed or

occurred for only 10 minutes. There is a possible world in which I exist for just those first few months, or even for just a few minutes; there is no possible world in which L exists for that period of time. Of course, if I had existed for, say, just ten minutes, there would have been an event which would have been my life, and which would have existed for just ten minutes; that event, however, would not have been L. We can put it like this: in any world in which I exist, there is an event which is my life; but it is not the case that there is an event which is my life, and which is my life in every world in which I exist.

(c), therefore, is false; it is not the case that I and the life of my body are such that necessarily, we exist at all the same times—that is, it is not the case that I and the life of my body are such that I have essentially the property of existing when and only when it does. But if (c) is false, the same goes for (b); since objections (2) and (3) both entail (b), both objections fail.

(4) If, as I say is possible, the replacement time for B and those parts is shorter than the assimilation time, there will be a brief period during which I don't have a body at all.<sup>15</sup> I will no longer have B, because all of B's parts have been replaced (and destroyed) during a time too brief for the new parts to be assimilated into B. I won't have any other body either, however; I won't have a body distinct from B, because there hasn't been time for these new parts to coalesce into a body. I therefore have no body at all during this time; there is no body that is my body at this time. How, then, can I continue to be conscious during this time, serenely reading the comics? Isn't it necessary that there be neurological activity supporting my consciousness during this time, if I am to be conscious then?

But is it logically necessary that there be neurological or other physical activity supporting my consciousness at any time at which I am conscious? That's a whopping

assumption. The most I need for my argument is that it is logically possible that I remain conscious during a brief period in which no neurological activity is supporting my consciousness; that's compatible with its being causally required that there be neurological activity when I am conscious. My entire argument has to do with what could happen; not with what would as a matter of fact happen, if this sort of replacement were to occur.<sup>16</sup> So the most that argument needs is that possibly, I exist and am conscious when no neurological activity is supporting my consciousness.<sup>17</sup> But the fact is it doesn't require even that. For consider a time  $t$  after the end of the replacement time but before the assimilation time has ended; let  $t$  be as close as you please to the end of the replacement time. At  $t$ , the replacing elements, the new parts, haven't yet had time to coalesce into a body. Nonetheless, any one of the new elements could be performing one of the several functions it will be performing when it has been integrated into a functioning human body. It could be playing part of the whole causal role it will be playing when the assimilation time has elapsed. In particular, therefore, the new neurons, before they have become part of a body, could be doing whatever it is they have to do in order to support consciousness. Accordingly, my argument requires that possibly I am conscious when I do not have a body; it does not require that possibly I am conscious when no neuronal or neurological activity is occurring.

## **II. Can a Material Thing Think? An Argument from Impossibility**

The Replacement Argument is an argument from possibility; as such, it proceeds from an intuition, the intuition that it is possible that my bodily parts, macroscopic or microscopic, be replaced while I remain conscious. But some people distrust modal intuitions. Of course it's impossible to do philosophy (or for that matter physics) without invoking modal intuitions of one sort or another or at any rate making modal declarations of one sort or another.<sup>18</sup> Still, it must be

conceded that intuition can sometimes be a bit of a frail reed. True, there is no way to conduct philosophy that isn't a frail reed, but intuition is certainly fallible. Further, some might think modal intuitions particularly fallible—although almost all of the intuitions involved in philosophy have important modal connections. Still further, one might think further that intuitions of possibility are especially suspect.<sup>19</sup> That is because it seems easy to confuse seeing the possibility of p with failing to see the impossibility of p. You can't see why numbers couldn't be sets; it doesn't follow that what you see is that they could be sets. Maybe I can't see why water couldn't be composed of something other than H<sub>2</sub>O; it doesn't follow that what I see is that water could be something other than H<sub>2</sub>O. And perhaps, so the claim might go, one who finds the replacement argument attractive is really confusing seeing the possibility of the replacements in question with failing to see their impossibility. Granted: I can't see that these replacements are impossible; it doesn't follow that what I see is that they are indeed possible.

To be aware of this possible source of error, however, is to be forewarned and thus forearmed. But for those who aren't mollified and continue to distrust possibility intuitions, I have another argument for dualism—one that depends on an intuition, not, this time, of possibility, but of impossibility. One who distrusts possibility intuitions may think more kindly of intuitions of impossibility—perhaps because she thinks that for the latter there isn't any obvious analogue of the possible confusion between failing to see that something is impossible with seeing that it is possible. Or rather, while there is an analogue—it would be confusing failure to see the possibility of p with seeing the impossibility of p—falling into that confusion seems less likely. In any event, the argument I'll now propose is for the conclusion that no material objects can think—i.e., reason and believe, entertain propositions, draw inferences, and the like. But of course I can think; therefore I am not a material object.

## A. Leibniz's Problem

I (and the same goes for you) am a certain kind of thing: a thing that can think. I believe many things; I also hope, fear, expect, anticipate many things. I desire certain states of affairs (desire that certain states of affairs be actual). I am capable of making decisions. I am capable of acting, and capable of acting on the basis of my beliefs and desires. I am conscious; and conscious of a rich, kaleidoscopic constellation of feeling, mental images, beliefs, and ways of being appeared to, some of which I enjoy and some of which I dislike. Naturally enough, therefore, I am not identical with any object that lacks any or all of these properties. What I propose to argue next is that some of these properties are such that no material object can have them. Again, others have offered similar arguments. In particular, many have seen a real problem for materialism in consciousness: it is extremely difficult to see how a material object could be conscious, could enjoy that vivid and varied constellation of feelings, mental images and ways of being appeared to. Others have argued that a material object can't make a decision (although of course we properly speak, in the loose and popular sense, of the chess playing computer as deciding which move to make next). These arguments seem to me to be cogent.<sup>20</sup> Here, however, I want to develop another argument of the same sort, another problem for materialism, a problem I believe is equally debilitating, and in fact fatal to materialism. Again, this problem is not a recent invention; you can find it or something like it in Plato. Leibniz, however, offers a famous and particularly forceful statement of it:

17. It must be confessed, moreover, that perception, and that which depends on it, are inexplicable by mechanical causes, that is by figures and motions. And supposing there were a machine so constructed as to think, feel and have perception, we could conceive of it as enlarged and yet preserving the same proportions, so that we might enter it as



into a mill. And this granted, we should only find on visiting it, pieces which push one against another, but never anything by which to explain a perception. This must be sought for, therefore, in the simple substance and not in the composite or in the machine.<sup>21</sup>

Now Leibniz uses the word 'perception' here; he's really thinking of mental life generally. His point, in this passage, is that mental life—perception, thought, decision—cannot arise by way of the mechanical interaction of parts. Consider a bicycle; like Leibniz's mill, it does what it does by virtue of the mechanical interaction of its parts. Stepping down on the pedals causes the front sprocket to turn, which causes the chain to move, which causes the rear sprocket to turn, which causes the back wheel to rotate. By virtue of these mechanical interactions, the bicycle does what it does, i.e., transports someone from one place to another. And of course machines generally—jet aircraft, refrigerators, computers, centrifuges—do their things and accomplish their functions in the same way. So Leibniz's claim, here, is that thinking can't arise in this way. A thing can't think by virtue of the mechanical interaction of its parts.

Leibniz is thinking of mechanical interactions—interactions involving pushes and pulls, gears and pulleys, chains and sprockets. But I think he would say the same of other interactions studied in physics, for example those involving gravity, electro-magnetism, and the strong and weak nuclear forces. Call these 'physical interactions'. Leibniz's claim is that thinking can't arise by virtue of physical interaction among objects or parts of objects. According to current science, electrons and quarks are simple, without parts.<sup>22</sup> Presumably neither can think—neither can adopt propositional attitudes; neither can believe, doubt, hope, want, or fear. But then a proton composed of quarks won't be able to think either, at least by way of physical relations between its component quarks, and the same will go for an atom composed of protons and electrons, a

molecule composed of atoms, a cell composed of molecules, and an organ (e.g., a brain), composed of cells. If electrons and quarks can't think, we won't find anything composed of them that can think by way of the physical interaction of its parts.

Leibniz is talking about thinking generally; suppose we narrow our focus to belief (although the same considerations apply to other propositional attitudes). What, first of all, would a belief be, from a materialist perspective? Suppose you are a materialist, and also think, as we ordinarily do, that there are such things as beliefs. For example, you hold the belief that Marcel Proust is more subtle than Louis L'Amour. What kind of a thing is this belief? Well, from a materialist perspective, it looks as if it would have to be something like a long-standing event or structure in your brain or nervous system. Presumably this event will involve many neurons related to each other in subtle and complex ways. There are plenty of neurons to go around: a normal human brain contains some 100 billion. These neurons, furthermore, are connected with other neurons at synapses; a single neuron can be involved in several thousand synapses, and there are some  $10^{15}$  synaptic connections. The total number of possible brain states, then, is absolutely enormous, vastly greater than the  $10^{80}$  electrons they say the universe contains. And the total number of possible neuronal events, while no doubt vastly smaller, is still enormous. Under certain conditions, groups of neurons involved in such an event fire, producing electrical impulses that can be transmitted (with appropriate modification and input from other structures) down the cables of neurons that constitute effector nerves to muscles or glands, causing, e.g., muscular contraction and thus behavior.

From the materialist's point of view, therefore, a belief will be a neuronal event or structure of this sort. But if this is what beliefs are, they will have two very different sorts of properties. On the one hand there will be electrochemical or neurophysiological properties ('NP

properties,' for short). Among these would be such properties as that of involving  $n$  neurons and  $n^*$  connections between neurons, properties that specify which neurons are connected with which others, what the rates of fire in the various parts of the event are, how these rates of fire change in response to changes in input, and so on. But if the event in question is really a belief, then in addition to those NP properties it will have another property as well: it will have to have a content. It will have to be the belief that  $p$ , for some proposition  $p$ . If this event is the belief that Proust is a more subtle writer than Louis L'Amour, then its content is the proposition Proust is more subtle than Louis L'Amour. My belief that naturalism is all the rage these days has as content the proposition Naturalism is all the rage these days. (That same proposition is the content of the German speaker's belief that naturalism is all the rage these days, even though she expresses this belief by uttering the German sentence 'Der Naturalismus ist dieser Tage ganz gross in Mode'; beliefs, unlike sentences, do not come in different languages.) It is in virtue of having a content, of course, that a belief is true or false: it is true if the proposition which is its content is true, and false otherwise. My belief that all men are mortal is true because the proposition which constitutes its content is true, but Hitler's belief that the Third Reich would last a thousand years was false, because the proposition that constituted its content was false.<sup>23</sup>

And now the difficulty for materialism is this: how does it happen, how can it be, that an assemblage of neurons, a group of material objects firing away has a content? How can that happen? More poignantly, what is it for such an event to have a content? What is it for this structured group of neurons, or the event of which they are a part, to be related, for example, to the proposition Cleveland is a beautiful city in such a way that the latter is its content? A single neuron (or quark, electron, atom or whatever) presumably isn't a belief and doesn't have content; but how can belief, content, arise from physical interaction among such material entities as

neurons? As Leibniz suggests, we can examine this neuronal event as carefully as we please; we can measure the number of neurons it contains, their connections, their rates of fire, the strength of the electrical impulses involved, the potential across the synapses—we can measure all this with as much precision as you could possibly desire; we can consider its electro-chemical, neurophysiological properties in the most exquisite detail; but nowhere, here, will we find so much as a hint of content. Indeed, none of this seems even vaguely relevant to its having content. None of this so much as slyly suggests that this bunch of neurons firing away is the belief that Proust is more subtle than Louis L'Amour, as opposed, e.g., to the belief that Louis L'Amour is the most widely published author from Jamestown, North Dakota. Indeed, nothing we find here will so much as slyly suggest that it has a content of any sort. Nothing here will so much as slyly suggest that it is about something, in the way a belief about horses is about horses.

The fact is, we can't see how it could have a content. It's not just that we don't know or can't see how it's done. When light strikes photoreceptor cells in the retina, there is an enormously complex cascade of electrical activity, resulting in an electrical signal to the brain. I have no idea how all that works; but of course I know it happens all the time. But the case under consideration is different. Here it's not merely that I don't know how physical interaction among neurons brings it about that an assemblage of them has content and is a belief. No, in this case, it seems upon reflection that such an event could not have content. It's a little like trying to understand what it would be for the number seven, e.g., to weigh 5 pounds, or for an elephant (or the unit set of an elephant) to be a proposition. (Pace the late (and great) David Lewis, according to whom the unit set of an elephant could be a proposition; in fact, on his view, there are uncountably many elephants the unit sets of which are propositions.) We can't see how that could happen; more exactly, what we can see is that it couldn't happen. A number just isn't the

sort of thing that can have weight; there is no way in which that number or any other number could weigh anything at all. The unit set of an elephant, let alone the elephant itself, can't be a proposition; it's not the right sort of thing. Similarly, we can see, I think, that physical activity among neurons can't constitute content. There they are, those neurons, clicking away, sending electrical impulses hither and yon. But what has this to do with content? How is content or aboutness supposed to arise from this neuronal activity? How can such a thing possibly be a belief? But then no neuronal event can as such have a content, can be about something, in the way in which my belief that the number seven is prime is about the number seven, or my belief that the oak tree in my backyard is without leaves is about that oak tree.

Here we must be very clear about an important distinction. Clearly there is such a thing as indication or indicator meaning.<sup>24</sup> Deer tracks in my backyard indicate that deer have run through it; smoke indicates fire; the height of the mercury column indicates the ambient temperature; buds on the trees indicate the coming of spring. We could speak here of 'natural signs': smoke is a natural sign of fire and the height of the mercury column is a natural sign of the temperature. When one event indicates or is a natural sign of another, there is ordinarily some sort of causal or nomic connection, or at least regular association, between them by virtue of which the first is reliably correlated with the second. Smoke is caused by fire, which is why it indicates fire; measles cause red spots on your face, which is why red spots on your face indicate measles; there is a causal connection between the height of the mercury column and the temperature, so that the latter indicates the former.

The nervous systems of organisms contain such indicators. A widely discussed example: when a frog sees a fly zooming by, the frog's brain (so it is thought) displays a certain pattern of neural firing; we could call such patterns 'fly detectors'. Another famous example: some

anaerobic marine bacteria have magnetosomes, tiny internal magnets. These function like compass needles, indicating magnetic north. The direction to magnetic north is downward; hence these bacteria, which can't flourish in the oxygen-rich surface water, move towards the more oxygen-free water at the bottom of the ocean.<sup>25</sup> Of course there are also indicators in human bodies. There are structures that respond in a regular way to blood temperature; they are part of a complex feedback system that maintains a more or less constant blood temperature by inducing (e.g.) shivering if the temperature is too low and sweating if it is too high. There are structures that monitor the amount of sugar in the blood and its sodium content. There are structures that respond in a regular way to light of a certain pattern striking the retina, to the amount of food in your stomach, to its progress through your digestive system, and so on. Presumably there are structures in the brain that are correlated with features of the environment; it is widely assumed that when you see a tree, there is a distinctive pattern of neural firing (or some other kind of structure) in your brain that is correlated with and caused by it.

Now we can, if we like, speak of 'content' here; it's a free country. We can say that the mercury column, on a given occasion, has a certain content: the state of affairs correlated with its having the height it has on that occasion. We could say, if we like, that those structures in the body that indicate blood pressure or temperature or saline content have a content on a given occasion: whatever it is that the structure indicates on that occasion. We could say, if we like, that the neural structure that is correlated with my looking at a tree has a content: its content, we could say, is what it indicates on that occasion. We can also, if we like, speak of information in these cases: the structure that registers my blood temperature, we can say, carries the information that my blood temperature is thus and so.

What is crucially important to see, however, is that this sort of content or information has nothing as such to do with belief, or belief content. There are those who—no doubt in the pursuit of greater generality—gloss over this distinction. Donald T. Cambell, for example, in arguing for the relevance of natural selection to epistemology, claims that “evolution—even in its biological aspects—is a knowledge process . . . .”<sup>26</sup> Commenting on Cambell’s claim, Franz Wuketits explains that

The claim is based on the idea that any living system is a “knowledge-gaining system.” This means that organisms accumulate information about certain properties of their environment. Hence life generally may be described as an information process, or, to put it more precisely, an information-increasing process.<sup>27</sup>

At any rate Wuketits has the grace to put ‘knowledge’ in scare quotes here. Knowledge requires belief; correlation, causal or otherwise, is not belief; information and content of this sort do not require belief. Neither the thermostat nor any of its components believes that the room temperature is thus and so. When the saline content of my blood is too low, neither I nor the structure correlated with that state of affairs (nor my blood) believes the saline content is less than it should be—or, indeed, anything else about the saline content. Indication, carrying information, is not belief; indicator content is not belief content, and these structures don’t have belief content just by virtue of having indicator content. And now the point here: I am not, of course, claiming that material structures can’t have indicator content; obviously they can. What I am claiming is that they can’t have belief content: no material structure can be a belief.

Here someone might object as follows. “You say we can’t see how a neural event can have content; but in fact we understand this perfectly well, and something similar happens all the time. For there is, after all, the computer analogy. A computer, of course, is a material object,

an assemblage of wires, switches, relays, and the like. Now suppose I am typing in a document. Take any particular sentence in the document: say the sentence 'Naturalism is all the rage these days'. That sentence is represented and stored on the computer's hard disk. We don't have to know in exactly what way it's stored (it's plusses and minuses, or a magnetic configuration, or something else; it doesn't matter). Now the sentence 'Naturalism is all the rage these days' expresses the proposition Naturalism is all the rage these days. That sentence, therefore, has the proposition Naturalism is all the rage these days as its content. But then consider the analogue of that sentence on the computer disk: doesn't it, too, express the same proposition as the sentence it represents? That bit of the computer disk with its plusses and minuses, therefore, has propositional content. But of course that bit of the computer disk is also (part of) a material object (as is any inscription of the sentence in question). Contrary to your claim, therefore, a material object can perfectly well have propositional content; indeed, it happens all the time. But if a computer disk or an inscription of a sentence can have a proposition as content, why can't an assemblage of neurons? Just as a magnetic pattern has as content the proposition Naturalism is all the rage these days, so too a pattern of neuronal firing can have that proposition as content. Your claim to the contrary is completely bogus and you should be ashamed of yourself." Thus far the objector.

If the sentence or the computer disk really did have content, then I guess the assemblage of neurons could too. But the fact is neither does—or rather, neither has the right kind of content: neither has original content; each has, at most, derived content. For how does it happen that the sentence has content? It's simply by virtue of the fact that we human beings treat that sentence in a certain way, use the sentence in a certain way, a way such that if a sentence is used in that way, then it expresses the proposition in question. Upon hearing that sentence, I think of,



grasp, apprehend the proposition Naturalism is all the rage these days. You can get me to grasp, entertain, and perhaps believe that proposition by uttering that sentence. How exactly all this works is complicated and not at all well understood; but the point is that the sentence has content only because of something we, we who are already thinkers, do with it. We could put this by saying that the sentence has secondary or derived content; it has content only because we, we creatures whose thoughts and beliefs already have content, treat it in a certain way. The same goes for the magnetic pattern on the computer disk; it represents or expresses that proposition because we assign that proposition to that configuration. But of course that isn't how it goes (given materialism) with that pattern of neural firing. That pattern doesn't get its content by way of being used in a certain way by some other creatures whose thoughts and beliefs already have content. If that pattern has content at all, then, according to materialism, it must have original or primary content. And what it is hard or impossible to see is how it could be that an assemblage of neurons (or a sentence, or a computer disk) could have original or primary content. To repeat: it isn't just that we can't see how it's done, in the way in which we can't see how the sleight of hand artist gets the pea to wind up under the middle shell. It is rather that we can see, to at least some degree, that it can't be done, just as we can see that an elephant can't be a proposition, and that the number 7 can't weigh seven pounds.

### **B. Parity?**

Peter van Inwagen agrees that it is hard indeed to see how physical interaction among material entities can produce thought: “. . . it seems to me that the notion of a physical thing that thinks is a mysterious notion, and that Leibniz's thought-experiment brings out this mystery very effectively.”<sup>28</sup>

Now I am taking this fact as a reason to reject materialism and hence as an argument for dualism. But of course it is a successful argument only if there is no similar difficulty for substance dualism itself. Van Inwagen believes there is a similar difficulty for dualism:

For it is thinking itself that is the source of the mystery of a thinking physical thing. The notion of a non-physical thing that thinks is, I would argue, equally mysterious. How any sort of thing could think is a mystery. It is just that it is a bit easier to see that thinking is a mystery when we suppose that the thing that does the thinking is physical, for we can form mental images of the operations of a physical thing and we can see that the physical interactions represented in these images—the only interactions that can be represented in these images—have no connection with thought or sensation, or none we are able to imagine, conceive or articulate. The only reason we do not readily find the notion of a non-physical thing that thinks equally mysterious is that we have no clear procedure for forming mental images of non-physical things (Loc. Cit.).

So dualism is no better off than materialism; they both have the same problem. But what precisely is this problem, according to van Inwagen? “. . . we can form mental images of the operations of a physical thing and we can see that the physical interactions represented in these images—the only interactions that can be represented in these images—have no connection with thought or sensation or none we are able to imagine, conceive or articulate.” As I understand van Inwagen here, he is saying that we can imagine physical interactions or changes in a physical thing; but we can see that the physical interactions represented in those images have no connection with thought. We can imagine neurons in the brain firing; we can imagine electrical impulses or perhaps clouds of electrons moving through parts of neurons, or whole chains of

neurons; we can imagine neural structures with rates of fire in certain parts of the structure changing in response to rates of fire elsewhere in or out of that structure: but we can see that these interactions have no connection with thought. Now I'm not quite sure whether or not I can imagine electrons, or their movements, or electrical impulses; but it does seem to me that I can see that electrical impulses and the motions of electrons, if indeed there are any such things, have nothing to do with thought.

Another way to put van Inwagen's point: no change we can imagine in a physical thing could be a mental change, i.e., could constitute thought or sensation, or a change in thought or sensation. But then we can't imagine a physical thing's thinking: i. e., we can't form a mental image of a physical thing thinking. And this suggests that the problem for materialism is that we can't form a mental image of a material thing thinking. But the same goes, says van Inwagen, for an immaterial thing: we also can't imagine or form a mental image of an immaterial thing thinking. Indeed, we can't form a mental image of any kind of thinking thing: "My point," he says, "is that nothing could possibly count as a mental image of a thinking thing" (177). Materialism and dualism, therefore, are so far on a par; there is nothing here to incline us to the latter rather than the former.

Thus far van Inwagen. The thought of a physical thing's thinking, he concedes, is mysterious; that is because we can't form a mental image of a physical thing's thinking. But the thought of an immaterial thing's thinking is equally mysterious; for we can't form a mental image of that either. This, however, seems to me to mislocate the problem for materialism. What inclines us to reject the idea of a physical thing's thinking is not just the fact that we can't form a mental image of a physical thing's thinking. There are plenty of things of which we can't form a mental image, where we're not in the least inclined to reject them as impossible. As Descartes

pointed out, I can't form a mental image of a chiliagon, a 1000-sided rectilinear plane figure (or at least an image that distinguishes it from a 100-sided rectilinear plane figure); that doesn't even suggest that there can't be any such thing. I can't form a mental image of the number 79's being prime: that doesn't incline me to believe that the number 79 could not be prime; as a matter of fact I know how to prove that it is prime. The fact is I can't form a mental image of the number 79 at all—or for that matter of any number; this doesn't incline me to think there aren't any numbers.

Or is all that a mistake? Is it really true that I can't form a mental image of the number 7, for example? Maybe I can form an image of the number 7; when I think of the number seven, sometimes there is a mental image present; it's as if one catches a quick glimpse of a sort of partial and fragmented numeral 7; we could say that I'm appeared to numeral-7ly. When I think of the actual world, I am sometimes presented with an image of the Greek letter alpha; when I think of the proposition All men are mortal I am sometimes presented with a sort of fleeting, fragmentary, partial image of the corresponding English sentence. Sets are nonphysical, but maybe I can imagine the pair set of Mic and Martha; when I try, it's like I catch a fleeting glimpse of curly brackets, enclosing indistinct images that don't look a whole lot like Mic and Martha. But is that really imagining the number 7, or the actual world, or the pair set of Mic and Martha? Here I'm of two minds. On the one hand, I'm inclined to think that this isn't imagining the number 7 at all, but instead imagining something connected with it, namely the numeral 7 (and the same for the actual world and the set of Mic and Martha). On the other hand I'm a bit favorably disposed to the idea that that's just how you imagine something like the number 7; you do it by imagining the numeral 7. (Just as you state a proposition by uttering a sentence or uttering certain sounds.) So I don't really know what to say. Can I or can't I

imagine nonphysical things like numbers, propositions, possible worlds, angels, God? I'm not sure.

What is clear, here, is this: if imagining the numeral 7 is sufficient for imagining the number 7, then imagining, forming mental images of, has nothing to do with possibility. For in this same way I can easily imagine impossibilities. I can imagine the proposition all men are mortal being red: first I just imagine the proposition, e.g., by forming a mental image of the sentence 'All men are mortal', and then I imagine this sentence as red. I think I can even imagine that elephant's being a proposition. David Kaplan once claimed he could imagine his refuting Gödel's Incompleteness Theorem: he imagined the Los Angeles Times carrying huge headlines: 'UCLA PROF REFUTES GÖDEL; ALL REPUTABLE EXPERTS AGREE'. In this loose sense, most anything can be imagined; but then the loose sense has little to do with what is or isn't possible. So really neither the loose nor the strong sense of 'imagining' (neither the weak nor the strong version of imagination) has much to do with possibility. There are many clearly possible things one can't imagine in the strong sense; in the weak sense, one can imagine many things that are clearly impossible.

What is it, then, that inclines me to think a proposition can't be red, or a horse, or an even number? The answer, I think, is that one can just see upon reflection that these things are impossible. I can't form a mental image of a proposition's having members; but that's not why I think no proposition has members; I also can't form a mental image of a set's having members. It's rather that one sees that a set is the sort of thing that (null set aside) has members, and a proposition is the sort of thing that cannot have members. It is the same with a physical thing's thinking. True, one can't imagine it. The reason for rejecting the idea, thinking it impossible, however, is not that one can't imagine it. It's rather that on reflection one can see that a physical

object just can't do that sort of thing. I grant that this isn't as clear and obvious, perhaps, as that a proposition can't be red; some impossibilities (necessities) are more clearly impossible (necessary) than others. But one can see it to at least a significant degree. Indeed, van Inwagen might be inclined to endorse this thought; elsewhere he says: ". . . Leibniz's thought experiment shows that when we carefully examine the idea of a material thing having sensuous properties, it seems to be an impossible idea."<sup>29</sup> But (and here is the important point) the same clearly doesn't go for an immaterial thing's thinking; we certainly can't see that no immaterial thing can think. (If we could, we'd have a quick and easy argument against the existence of God: no immaterial thing can think; if there were such a person as God, he would be both immaterial and a thinker; therefore ...)

Van Inwagen has a second suggestion:

In general, to attempt to explain how an underlying reality generates some phenomenon is to construct a representation of the working of that underlying reality, a representation that in some sense "shows how" the underlying reality generates the phenomenon. Essentially the same considerations as those that show that we are unable to form a mental image that displays the generation of thought and sensation by the workings of some underlying reality (whether the underlying reality involves one thing or many, and whether the things it involves are physical or non-physical) show that we are unable to form any sort of representation that displays the generation of thought and sensation by the workings of an underlying reality. (Metaphysics, pp. 177-78)

The suggestion is that we can't form an image or any other representation displaying the generation of thought by way of the workings of an underlying reality; hence we can't see how it

can be generated by physical interaction among material objects such as neurons. This much seems right—at any rate we certainly can't see how thought could be generated in that way. Van Inwagen goes on to say, however, that this doesn't favor dualism over materialism, because we also can't see how thought can be generated by the workings of an underlying non-physical reality. And perhaps this last is also right. But here there is an important dissimilarity between dualism and materialism. The materialist thinks of thought as generated by the workings of an underlying reality—i.e., by the physical interaction of such physical things as neurons; the dualist, however, typically thinks of an immaterial self, a soul, a thing that thinks, as simple. An immaterial self doesn't have any parts; hence, of course, thought isn't generated by the interaction of its parts. Say that a property P is basic to a thing x if x has P, but x's having P is not generated by the interaction of its parts. Thought is then a basic property of selves, or better, a basic activity of selves. It's not that (for example) there are various underlying immaterial parts of a self whose interaction produces thought. Of course a self stands in causal relation to its body: retinal stimulation causes a certain sort of brain activity which (so we think) in turn somehow causes a certain kind of experience in the self. But there isn't any way in which the self produces a thought; it does so immediately. To ask "How does a self produce thought?" is to ask an improper question. There isn't any how about it.

By way of analogy: consider the lowly electron. According to current science, electrons are simple, not composed of other things. Now an electron has basic properties, such as having a negative charge. But the question 'How does an electron manage to have a charge?' is an improper question. There's no how to it; it doesn't do something else that results in its having such a charge, and it doesn't have parts by virtue of whose interaction it has such a charge. Its having a negative charge is rather a basic and immediate property of the thing (if thing it is). The

same is true of a self and thinking: it's not done by underlying activity or workings; it's a basic and immediate activity of the self. But then the important difference, here, between materialism and immaterialism is that if a material thing managed to think, it would have to be by way of the activity of its parts: and it seems upon reflection that this can't happen.<sup>30</sup> Not so for an immaterial self. Its activity of thinking is basic and immediate. And it's not the case that we are inclined upon reflection to think this can't happen—there's nothing at all against it, just as there is nothing against an electron's having a negative charge, not by virtue of the interaction of parts, but in that basic and immediate way. The fact of the matter then is that we can't see how a material object can think—that is, upon reflection it seems that a material object can't think. Again, not so for an immaterial self.

True, as van Inwagen says, thought can sometimes seem mysterious and wonderful, something at which to marvel. (Although from another point of view it is more familiar than hands and feet). But there is nothing here to suggest that it can't be done. I find myself perceiving my computer; there is nothing at all, here, to suggest impossibility or paradox. Part of the mystery of thought is that it is wholly unlike what material objects can do: but of course that's not to suggest that it can't be done at all. Propositions are also mysterious and have wonderful properties: they manage to be about things; they are true or false; they can be believed; they stand in logical relations to each other. How do they manage to do those things? Well, certainly not by way of interaction among material parts. Sets manage, somehow, to have members—how do they do a thing like that? And why is it that a given set has just the members it has? How does the unit set of Lance Armstrong manage to have just him as a member? What mysterious force, or fence, keeps Leopold out of that set? Well, it's just the nature of sets to be like this. These properties can't be explained by way of physical interactions among material



parts, but that's nothing at all against sets. Indeed, these properties can't be explained at all. Of course if you began with the idea that everything has to be a material object, then thought (and propositions and sets) would indeed be mysterious and paradoxical. But why begin with that idea? Thought is seriously mysterious, I think, only when we assume that it would have to be generated in some physical way, by physical interaction among physical objects. That is certainly mysterious; indeed it goes far beyond mystery, all the way to apparent impossibility. But that's not a problem for thought; it's a problem for materialism.

### **III. Arguments for Materialism**

The above arguments for dualism and others like them are powerful. Like philosophical arguments generally, however, they are not of that wholly apodictic and irrefragable character Kant liked to claim for his arguments; they are defeasible. It is possible to disregard or downgrade the intuitions of possibility and impossibility to which they appeal. Further, if there were really powerful arguments for materialism—stronger than these arguments against it—then perhaps the appropriate course would be to embrace materialism. But are there any such powerful arguments?

No—or at least I've never seen any. There is the old chestnut according to which no immaterial object can cause changes in the hard, heavy, massive, massy (messy) physical world; there is the claim that dualism, or at least interactionistic dualism, violates the principle of Conservation of Energy; there is the charge that dualism is unscientific; there is the complaint that soul stuff is hard to understand; there is the canard that dualism is explanatorily impotent. None of these has any force at all.<sup>31</sup> However there is one that is perhaps not completely without promise. According to Nancey Murphy:

In particular, nearly all of the human capacities or faculties once attributed to the soul are now seen to be functions of the brain. Localization studies—that is, finding regional structures or distributed systems in the brain responsible for such things as language, emotion and decision making—provide especially strong motivation for saying that it is the brain that is responsible for these capacities, not some immaterial entity associated with the body. In Owen Flanagan’s terms, it is the brain that is the res cogitans—the thinking thing.<sup>32</sup>

Localization studies show that when certain kinds of mental activity occur, certain parts of the brain display increased blood flow and increased electrical activity. Paul Churchland goes on to point out that mental activity is also in a certain important way dependent on brain activity and brain condition:

Alcohol, narcotics, or senile degeneration of nerve tissue will impair, cripple, or even destroy one’s capacity for rational thought. Psychiatry knows of hundreds of emotion-controlling chemicals (lithium, chlorpromazine, amphetamine, cocaine, and so on) that do their work when vectored into the brain. And the vulnerability of consciousness to the anesthetics, to caffeine, and to something as simple as a sharp blow to the head, shows its very close dependence on neural activity in the brain. All of this makes perfect sense if reason, emotion and consciousness are activities of the brain itself. But it makes very little sense if they are activities of something else. We may call this the argument from the neural dependence of all known mental phenomena.<sup>33</sup>

Of course it isn’t true that it makes very little sense to say that activities of the immaterial self or soul are dependent in this way on the proper function of the brain; still, this argument

from localization and neural dependence is perhaps the strongest of the arguments against dualism. That may not be much of a distinction; the other arguments, I believe, are without any force at all. But perhaps this argument has a little something to be said for it; at any rate dependence and localization phenomena do suggest the possibility that the brain is all there is. Taken as an argument, however, and looked at in the cold light of morning, it has little to be said for it. What we know, here, is that for at least many mental functions or actions M, there are parts of the brain B such that (1) when M occurs, there is increased blood flow and electrical activity in B, and (2) when B is damaged or destroyed, M is inhibited or altogether absent. Consider, therefore, the mental activity of adding a column of figures, and let's assume that there is a particular area of the brain related to this activity in the way suggested by (1) and (2). Does this show or tend to show that this mental activity is really an activity of the brain, rather than of something distinct from the brain?

Hardly. There are many activities that stand in that same or similar relation to the brain. Consider walking, or running, or speaking, or waving your arms or moving your fingers: for each of these activities too there is a part of your brain related to it in such a way that when you engage in that activity, there is increased blood flow in that part; and when that part is damaged or destroyed, paralysis results so that you can no longer engage in the activity. Who would conclude that these activities are really activities of the brain rather than of legs and trunk, or mouth and vocal cords, or arms? Who would conclude that your fingers' moving is really an activity of your brain and not of your fingers? Your fingers' moving is dependent on appropriate brain activity; it hardly follows that their moving just is an activity of your brain. Digestion will occur only if your brain is in the right condition; how does it follow that digestion is really an activity of the brain, and not an activity of the digestive system? Your brain's functioning

properly depends on blood flow and on the proper performance of your lungs; shall we conclude that brain function is really circulatory or pulmonary activity? All of your activities depend upon your ingesting enough and the right kind of food; shall we see here vindication of the old saw ‘you are what you eat’? The point, obviously, is that dependence is one thing, identity quite another. Appropriate brain activity is a necessary condition for mental activity; it simply doesn’t follow that the latter just is the former. Nor, as far as I can see, is it even rendered probable. We know of all sorts of cases of activities A that depend upon activities B but are not identical with them. Why should we think differently in this case?

Perhaps a more promising way of developing this argument would go as follows. In science, it is common to propose identities of various kinds: water is identical with H<sub>2</sub>O, heat and pressure with molecular motion, liquidity, solidity, gaseousness with certain properties of assemblages of molecules, and so on. This kind of identification, it might be argued, is theoretically useful in at least two ways; in some cases it provides explanations, answers to questions that are otherwise extremely difficult to answer, and in others finesses the questions by obviating the need for answers, showing instead that the question itself is bogus, or ill-formed, or has a wholly trivial answer. Well, why not the same here? Suppose we identify mental activity with brain activity; more precisely, suppose we identify such properties as being in pain and being conscious with such properties as having C-fibers that are firing and displaying activity in the pyramidal cells of layer 5 of the cortex involving reverberatory circuits.<sup>34</sup> Then first of all, we don’t have to answer the otherwise difficult questions, “Why is it that when someone is in pain, the C-fibers in her brain are firing?” Or “Why is it that when someone is conscious, his brain is displaying activity in the pyramidal cells?” (Alternatively, they might have answers, but the answer would be pretty easy: “Because being in pain just is having firing C-fibers and being

conscious just is displaying pyramidal activity.) And second we will be able to answer some questions otherwise very difficult: for example, “Why is it that rapping someone smartly over the head interferes with their ability to follow a proof of Gödel’s Theorem?” This identification, therefore, is theoretically fruitful and hence justified by the principle of inference to the best explanation (B&S, pp. 24, 45). Still further, of course, if mental properties are really identical with and thus reduced to neurophysiological properties of the brain, then dualism will be false. So here we have another objection to dualism.

Now first, note the language involved here: the suggestion is that we identify, say, the property of being in pain with the property of having firing C-fibers. That makes it sound as if it’s just up to us whether these properties are identical—we can just identify them, if we find that useful. But of course it isn’t just up to us, and we can’t really do any such thing. All we can do is declare, perhaps loudly and slowly, that these properties are identical; but saying so doesn’t make it so (not even if your peers let you get away with so saying).

More important, what about the fact that these properties—being in pain and having firing C-fibers, for example, or being conscious and displaying activity in the pyramidal cells seem so utterly different? Pain and consciousness are immediately apprehended phenomenal properties; not so for firing C-fibers or active pyramidal cells. And as for that pyramidal activity, if that’s what being conscious just is, then nothing, not even God, could be conscious but not have those pyramidal cells. So do we have here another shiny new argument for atheism, this time from neural science: God, if he exists, is conscious, but without a body; neuroscience shows that being conscious just is having active pyramidal cells; hence . . . ? On the face of it, these properties seem at least as different as being chalk and being cheese. In fact on the face of it they seem more different than the latter; at least any pair of things that exemplify being chalk

and being cheese are clearly both material objects. How can being in pain be the same property as having firing C- fibers when it seems so utterly clear that someone could be in pain without having C-fibers that are firing, as well as have firing C-fibers without being in pain? Perhaps it's true that if these properties were identical, we would have answers to some otherwise difficult questions (and avoid some other questions): but isn't it obvious that the properties are not identical? Are not both

(1) Possibly, someone is in pain when no C-fibers are firing

and

(2) Possibly, C-fibers are firing when no one is in pain

wholly obvious?

Well, they certainly look obvious. Maybe identifying these properties would have a theoretical payoff; but the properties just don't seem to be identical. You might as well 'identify' Bill with his essence: concrete objects are so unruly and messy, after all. And why stop with Bill? Why not identify every concrete object with its essence, thus finessing all those annoying questions about the relation between concrete objects and abstract properties? There is a problem about how God knows future contingents: how does he know that tomorrow I will freely go for a bike ride? It hasn't happened yet, and since it will be a free action when it does happen, he can't deduce it from present conditions and causal laws. No problem, mates; just identify truth with the property being believed by God. How can we so blithely declare these properties identical when they look so different? How can we declare (1) and (2) false when they seem so obviously true?

Now here appeal will be made to Kripke and his celebrated thesis about necessary but a posteriori propositions. B&S and others suggest that the appearance of falsehood for (1) and (2) is like the appearance of contingency of such propositions as

(3) Water is H<sub>2</sub>O,

or

(4) Gold has atomic number 79.

In these cases we have the appearance of contingency; but, so the claim goes, the appearance is shown by Kripkean considerations to be illusory. We initially think that these propositions are contingent; Kripke shows us that in fact they are necessary. As Sydney Shoemaker says in a similar context,

Kripke . . . argued that the class of truths deserving this label [i.e., the label of being necessary] is much larger than had traditionally been supposed. And, in his most radical departure from the traditional view, he held that many of these truths have the epistemic status of being a posteriori.<sup>35</sup>

Among these truths, of course, are (3) and (4). But then once we see that this is how it goes in the case of water and H<sub>2</sub>O, and being gold and having atomic number 79, we can apply the lesson to neurophysiological and mental properties. Indeed, according to B&S,

The crucial question for the issue we have been discussing in this paper is whether a relevant contrast can be shown between the relation between water and H<sub>2</sub>O on the one hand and the relation between consciousness and some brain process on the other (43).

But such a relevant contrast, I believe, can easily be shown. Suppose we look a bit more deeply into the relevant Kripkean considerations. Kripke's principal thesis here, of course, is that

natural kind terms--'tiger', 'water', 'gold'—function as rigid designators; they are not, for example, as Frege or Russell thought, disguised or abbreviated definite descriptions.<sup>36</sup> While there is a certain amount of controversy about the notion of rigid designation, what is clear is that the thesis in question is a semantical thesis, a thesis about the meaning or function of certain terms. And that should put us on our guard. A semantical thesis about how certain terms work is not, just by itself, of direct relevance to the modal question which propositions are necessary or contingent; what it is relevant to, is the question which propositions get expressed by which sentences. Pace Shoemaker, a semantical thesis can't by itself show us that the class of necessary propositions is larger than we thought; what it can show us is that sentences we thought expressed contingent propositions really express necessary propositions.<sup>37</sup>

Accordingly, consider the sentences

(5) 'Water is H<sub>2</sub>O'

and

(6) 'Gold has atomic number 79';

what Kripke shows us is that these sentences, contrary to what we perhaps originally thought, really express necessary rather than contingent propositions. But it wasn't that we were clear about which propositions were in fact expressed by those sentences, and Kripke got us to see that those propositions, contrary to what we thought, were necessary. It is rather that he corrected our ideas about which propositions are expressed by those sentences: we mistakenly thought (5) expressed a certain proposition P which we correctly thought to be contingent; in fact (5) expresses a different proposition Q, a proposition that appears to be necessary. We might have thought that 'water' is synonymous with something like "the clear tasteless odorless stuff we find



in lakes and streams,” in which case the sentence (5) expresses a proposition put more explicitly by

(7) The clear, tasteless odorless stuff found in lakes and streams is H<sub>2</sub>O.

This is clearly contingent: it entails the contingent proposition that H<sub>2</sub>O is found in lakes and streams. By way of a judicious selection of examples, however, Kripke gets us to see (if he's right) that the proposition expressed by (5) isn't (7) at all. What proposition does it express? That proposition can be put as follows.

(8) Consider the stuff actually to be found in the rivers and lakes: that stuff is H<sub>2</sub>O.

That's the proposition expressed by (5). Alternatively, consider the stuff we do in fact find in lakes and rivers and name it 'XX'; then

(8) XX is H<sub>2</sub>O.

(8) is at least arguably necessary:<sup>38</sup> it seems sensible to think that very stuff, i.e., H<sub>2</sub>O, could not have failed to be H<sub>2</sub>O. We are inclined to think, perhaps under the influence of mistaken views about the function of kind terms, that 'water' expresses such properties as being clear, tasteless, odorless and filling the lakes and streams. Kripke gets us to see that 'water' does not express those properties, which could be had by very many different substances, but is instead a rigid designator of the stuff that actually has those properties, i.e., as I would put it, expresses the (or an) essence of that stuff. The difference is between, on the one hand, the term's expressing the properties we use to fix its reference, and, on the other, the term's being a rigid designator (I'd say expressing the essence) of what it denotes when its reference is fixed in that way. By analogy, return to those thrilling days of yesteryear, when Quine asked us to consider such sentences as

(9) Hesperus is identical with Venus.

This may look contingent: we might think ‘Hesperus’ expresses the property of being the evening star, i.e., of being the first heavenly body to appear in the evening. Surely it’s not necessary that the first heavenly body to appear in the evening is Venus—any number of other heavenly bodies could have been (and I guess sometimes actually are) the first to appear in the evening. But what Kripke got us to see is that in fact ‘Hesperus’ does not express that property; it is instead a name or rigid designator (expresses an essence of) the thing that has that property, in which case (9), contrary to what we might have thought, does not express a contingent proposition after all.

How does this apply to the case in question, the case of the proposed identification of mental properties with neurophysiological properties? As follows: in the water/ H<sub>2</sub>O case, what we learn from Kripke is not that some proposition we had thought contingent is really necessary; what we learn instead is that some sentence we thought expressed a contingent proposition really expresses a necessary proposition. What we learn is a semantical fact, not a modal fact. It isn’t that there is some proposition we thought to be contingent and is now seen to be necessary. It isn’t that (3), the proposition, formerly appeared to us to be contingent, but then was seen, via Kripkean considerations, to be necessary; it is rather that (5) formerly seemed to express a contingent proposition and is now seen to express a necessary proposition. Our problem was not modal illusion, but semantical illusion. So what we have is not a reason for mistrusting modal intuition (more specifically, an intuition of possibility); it is rather a reason (if only a weak one) for mistrusting our ideas about the semantics of proper names and kind names.

But then there is a large, important and crucial difference between the water/ H<sub>2</sub>O case and the pain/firing C-fibers case. In the former, as I’ve just been arguing, the proposed identification doesn’t conflict with any modal intuitions at all. Indeed, (3) does seem intuitively

to be necessary: that very stuff could not have been something other than H<sub>2</sub>O. So here there isn't so much as a hint of conflict with modal intuition. In the latter case however, the case of pain/C-fibers firing, there is a clear and wholly obvious conflict with intuition.

(10) Someone is in pain when no C-fibers are firing;

appears for all the world to be possible; similarly, of course, for

(11) Someone is conscious when there is no pyramidal cell activity.

According to the proposed identification, however, these propositions are (of course) impossible.

If pain just is the firing of C-fibers and consciousness just is pyramidal cell activity, then (10)

and (11) are equivalent, in the broadly logical sense, to

(12) C-fibers are firing when no C-fibers are firing

and

(13) There is pyramidal cell activity when there is no pyramidal cell activity.

In other words, the proposed identification of water with H<sub>2</sub>O goes contrary to no modal

intuition; the proposed identification of pain with C-fiber firing, and consciousness with

pyramidal cell activity, on the other hand, is wholly counter-intuitive. The latter identifications

go directly against strong modal intuitions; the former does not.

This objection to dualism, therefore, is no stronger than the others. No doubt splendid theoretical advantages would be forthcoming from the identification of mental with

neurophysiological properties, as with the identification of concrete objects with their essences.

But these theoretical advantages are surely outweighed by the fact that the proposed

identifications are obviously false. Like the other objections to dualism, accordingly, this one is

without any force. In conclusion, then: there are powerful arguments against materialism and

none for it. Why, therefore, should anyone want to be a materialist?<sup>39</sup>

---

<sup>1</sup> See, e.g., Richard Swinburne, The Evolution of the Soul (Oxford: Clarendon Press, 1987), pp. 2, 145. Aquinas' position on the relation between soul and body may be a special case of this view; see my "Materialism and Christian Belief," in Persons: Human and Divine, ed. Peter van Inwagen and Dean Zimmerman, forthcoming.

<sup>2</sup> See, e.g., Bas van Fraassen, The Empirical Stance (New Haven: Yale, 2002), pp. 50ff.

<sup>3</sup> Although I can't help concurring with David Armstrong, no friend of dualism:

But disembodied existence seems to be a perfectly intelligible supposition. . . . Consider the case where I am lying in bed at night thinking. Surely it is logically possible that I might be having just the same experiences and yet not have a body at all. No doubt I am having certain somatic, that is to say, bodily sensations. But if I am lying still these will not be very detailed in nature, and I can see nothing self-contradictory in supposing that they do not correspond to anything in physical reality. Yet I need be in no doubt about my identity. (A Materialist Theory of Mind [London: Routledge, 1968], p. 19.)

<sup>4</sup> See my Warrant and Proper Function (New York; Oxford University Press, 1993), chap. 6.

<sup>5</sup> See, e.g., Dean Zimmerman, "Material People" in The Oxford Handbook of Metaphysics (Oxford: Clarendon Press, 2002), pp. 504ff. Zimmerman himself seems attracted to the thought that "the mass of matter" of which one's body is composed is an object distinct from the latter, but colocated with it (although of course he is not attracted to the idea that a person is identical with such a mass of matter).

<sup>6</sup> See Zimmerman, loc cit.

<sup>7</sup> One such argument, for example, apparently has the following form:

(a) Many people who advocate p, do so in the service of a hope that science will never be able to explain p;

therefore

(b) not-p.

See Daniel Dennett, Darwin's Dangerous Idea (New York: Simon and Schuster, 1995), p. 27.

Another seems to have the form

(a) If you believe p, prestigious people will laugh at you;

therefore

(b) not-p

(or perhaps

---

(b\*) don't believe p?)

See Daniel Dennett, Explaining Consciousness (Boston: Little, Brown and Co., 1991), p. 37.

<sup>8</sup> “But on the kinds of figures that are coming out now, it seems like the whole brain must get recycled about every other month.” John McCrone, “How Do You Persist When Your Molecules Don't?” Science and Consciousness Review, (web-journal, June 2004, No. 1).

<sup>9</sup> Here I am indebted especially to Michael Rea.

<sup>10</sup> See, e.g., David Hershenov, “The Metaphysical Problem of Intermittent Existence and the Possibility of Resurrection, Faith and Philosophy (Jan. 2003), p. 33.

<sup>11</sup> Complaint: this new ‘part’ as you call it, isn't really a part, at first, anyway, because at first it isn't yet integrated into the organism. Reply: think of ‘part’ here, as like ‘part’ in ‘auto parts store’. Would you complain that the auto parts store is guilty of false advertising, on the grounds that none of those carburetors, spark plugs and piston rings they sell is actually part of an automobile?

<sup>12</sup> Here I'm indebted to Richard Fumerton.

<sup>13</sup> No doubt this flexibility results from a principle of compounding for events: given any two successive events  $e_1$  and  $e_2$  occurring at roughly the same place, there is another event  $e_3$  compounded of them, an event that has each of them as a subevent. Short of such metaphysical extravaganzas as mereological universalism, clearly there is no corresponding principle for material objects. I stand in the corner from  $t_1$  to  $t_2$ ; then I leave and you stand in the corner from  $t_2$  to  $t_3$ ; it doesn't follow that there is a material object that stands in the corner from  $t_1$  to  $t_3$  and has my body and your body as successive parts.

<sup>14</sup> See Dean Zimmerman, op. cit., p. 510.

<sup>15</sup> Here I am indebted to Nicholas Wolterstorff.

<sup>16</sup> And hence strictly speaking, the argument doesn't require a thought experiment; it requires instead seeing that a certain state of affairs or proposition is possible. See George Bealer, “Intuition and the Autonomy of Philosophy” in Rethinking Intuition, ed. Michael DePaul and William Ramsey (New York: Rowman and Littlefield, 1998), p. 207.

<sup>17</sup> Not strictly relevant, but of interest: could I perhaps be a computer (hardware), a computer made of flesh and blood? There are three possibilities here: I might be the hardware, I might be the program, and I might be the mereological sum of the hardware and the program. The first suggestion is vulnerable to the macroscopic Replacement Argument; on the other two, I would not be a material object. So no help for materialism there.

<sup>18</sup> Realists will say that there can't be similarity without a property had by the similar things, thus resting on an alleged intuition of impossibility; nominalists will deny this claim, thus resting on an alleged intuition of possibility. In his argument for indeterminacy of translation, Quine claims that the native's behavior is consistent with his meaning ‘rabbit state’ or ‘undetached rabbit part’ or ‘rabbit’ by ‘gavagai’, thus (despite his animadversions) relying on an intuition of possibility. Similarly for his and

others' claims about the underdetermination of theory by evidence. Further, anyone who proposes an analysis (of knowledge, for example) relies on intuition, as does someone who objects to such an analysis (by proposing a Gettier case, for example). In philosophy of mind we have Jackson's Mary example, Burge's arthritis example, twin earth arguments for a posteriori necessities and wide content, refutations of phenomenalism and behaviorism, and much else besides, all of which rely centrally and crucially on modal intuition. Most arguments for materialism rely on modal intuition (for example the intuition that an immaterial thing can't cause effects in the hard, heavy, massive material world). Indeed, take your favorite argument for any philosophical position: it will doubtless rely on modal intuition.

<sup>19</sup> See below, pp. 46ff.

<sup>20</sup> There is also the complex but powerful argument offered by Dean Zimmerman, op. cit., pp. 517ff..

<sup>21</sup> Monadology 17. In Leibniz Selections, ed. Philip Weiner (New York: Charles Scribner's Sons, 1951), p. 536.

<sup>22</sup> Although there are speculative suggestions that quarks may in fact be composed of strings.

<sup>23</sup> I've been assuming that there really are such things as beliefs. A materialist might demur, taking a leaf from those who accept 'adverbial' accounts of sensation, according to which there aren't any red sensations or red sense data or red appearances: what there are instead are cases of someone's sensing redly or being appeared to redly. Similarly, the materialist might claim that there isn't any such thing as the belief that all men are mortal (or any other beliefs); what there is instead are cases of people who believe in the all-men-are-mortal way. This may or may not make sense; if it does make sense, however, a person will presumably believe in the all-men-are-mortal way only if she harbors a neuronal structure or event that has as content the proposition all men are mortal.

<sup>24</sup> See Fred Dretske's Explaining Behavior (Cambridge, MA: MIT Press, 1988), pp. 54ff. See also Bill Ramsey's Using and Abusing Representation: Reassessing the cognitive Revolution (presently unpublished). Materialists who try to explain how a material structure like a neuronal event can be a belief ordinarily try to do so by promoting indicators to beliefs; for animadversions on such attempts, see the appendix in my "Materialism and Christian Belief."

<sup>25</sup> Dretske, op. cit., p. 63.

<sup>26</sup> "Evolutionary Epistemology" in The Philosophy of Karl Popper, ed. P.A. Schilpp (LaSalle: Open Court, 1974), p. 413.

<sup>27</sup> "Evolutionary Epistemology" in Biology and Philosophy, vol. 1, No. 2 (1986), p. 193.

<sup>28</sup> Metaphysics (Boulder, Colorado, 2002 (second edition)), p. 176.

<sup>29</sup> "Dualism and Materialism: Athens and Jerusalem?", Faith and Philosophy 12:4, Oct. 1995, p. 478. That is (I take it), it seems to be necessary that material things don't have such properties. Van Inwagen's examples are such properties as being in pain and sensing redly; the same goes, I say, for properties like being the belief that p for a proposition p.

<sup>30</sup> But couldn't a material thing also just directly think, without depending on the interaction of its parts? According to Pierre Cabanis, "The brain secretes thought as the liver secretes bile;" couldn't we think of this as the brain (or, if you like, the whole organism) directly thinking, not by way of the interaction of its parts? Well, if that's how a brain thinks, it isn't like the way a liver secretes bile; the latter certainly involves the liver's having parts, and those parts working together in the appropriate way. Further, the idea of a physical thing's thinking without the involvement of its parts is even more clearly impossible than that of a physical thing's thinking by virtue of the interaction of its parts. Aren't those neurons in the brain supposed to be what enables it to think? You might as well say that a tree or my left foot thinks. Consider any nonelementary physical object—a tree, an automobile, perhaps a horse: such a thing does what it does by virtue of the nature and interaction of its parts. Are we to suppose that some physical object—a brain, let's say—does something like thinking apart from involvement of its parts? Talk about appealing to magic!

<sup>31</sup> This may seem a bit abrupt; for substantiation, see my "Materialism and Christian Belief," (footnote 1)

<sup>32</sup> Brown, Murphy and Malony, Whatever Happened to the Soul? (Minneapolis: Fortress Press, 1998), p. 1.

<sup>33</sup> Matter and Consciousness (Cambridge: MIT Press, 1984), p 20. See also Thomas Nagel's "The Psychophysical Nexus," in Concealment and Exposure and Other Essays (New York: Oxford University Press, 2002); in the course of a long, detailed and subtle discussion, Thomas Nagel argues that there is a logically necessary connection between mental states and physical states of the following sort: for any mental state M there is a physical state P such that there is some underlying reality R, neither mental nor physical but capable of having both mental and physical states, which has essentially the property of being such that necessarily, it is in P just if it is in M. (And perhaps it would be sensible to go on from that claim to the conclusion that it is not possible that I exist when my body B does not.) Nagel concedes that it seems impossible that there be such a reality; his argument that nonetheless there really is or must be such a thing is, essentially, just an appeal to localization/dependency phenomena: "The evident massive and detailed dependence of what happens in the mind on what happens in the brain provides, in my view, strong evidence that the relation is not contingent but necessary (p 202), and "The causal facts are strong evidence that mental events have physical properties, if only we could make sense of the idea" (p. 204). The particular route of his argument here is via an argument to the best explanation: he suggests that the only really satisfactory explanation of those localization/dependency phenomena is the existence of such an underlying reality. (Of course if that is what it takes for a really satisfying explanation, it is less than obvious that there is a really satisfying explanation here.) This argument has also made its way into the popular press: See Steven Pinker's "How to Think About the Mind," Newsweek (Sept. 27, 2004), p. 78.

<sup>34</sup> See Ned Block and Robert Stalnaker, "Conceptual Analysis, Dualism, and the Explanatory Gap" (hereafter 'B&S'), Philosophical Review, Jan. 1999, p. 1.

<sup>35</sup> "Causal and Metaphysical Necessity," Pacific Philosophical Quarterly 79 (1998), p. 59.

<sup>36</sup> Or if they are disguised descriptions, the descriptions they disguise express essences of their donotata.

---

<sup>37</sup>Not everyone is prepared to distinguish propositions from sentences. Those who do not make that distinction, however, will presumably be able to make an equivalent distinction by noting the difference between coming to see that a sentence is necessary in virtue of discovering that it doesn't mean what one thought it did, and coming to see that it is necessary without learning anything new about its meaning.

<sup>38</sup>For the moment ignore the fact that it's contingent that there is any such thing as H<sub>2</sub>O.

<sup>39</sup>In addition to the people mentioned in the text, I thank Michael Bergmann, Evan Fales, Trenton Merricks, William Ramsey, and the members of the Notre Dame Center for Philosophy of Religion discussion group, in particular Thomas Flint and Peter van Inwagen, as well as others I have inadvertently overlooked. I'm especially grateful to Dean Zimmerman.