

Editorial

Fodor's Gloom,

or

What does it mean that dualism seems true?

Eric Dietrich

Program in Philosophy, Computers, & Cognitive Science

Binghamton Univ., Binghamton, NY

1. Difficult problems and things we can't know.

Any time you have philosophers working on a problem, you know you've got troubles. If a question has attracted the attention of the philosophers that means that either it is intractably difficult with convolutions and labyrinthine difficulties that would make other researchers blanch, or that it is just flat out impossible to solve. Impossible problems masquerade as intractable problems until someone either proves the problem is impossible (which can only happen in mathematics), or someone shows all solutions to the problem violate laws of physics (like the perpetual motion machine, for example), or until enough people fail so that declaring defeat is a reasonable move. The problem of consciousness is prototypical of this latter case. Indeed, one might say that it is the Platonic ideal of such a problem. The mere fact that philosophers wrestle with the problem of consciousness should be regarded by psychologists of all stripes as extremely bad news. If the philosophers can't make any headway, psychologists are doomed.

Interestingly, in the early part of this century, the philosophers got their act together, after a fashion, and declared that all the problems they had been working on for centuries -- the problem of freewill, the question of what principles ought to govern how we treat each other, the problem of other minds, the puzzle about the nature of causation and the nature of scientific laws, and the question about what constitutes art, (to name a few of the more famous problems) -- were in fact impossible to solve, and impossible for a very peculiar reason: all philosophical problems, declared the early twentieth century philosophers, were really in fact *pseudoproblems*. Pseudoproblems are things that seem like real problems, but are in fact not problems at all, and are actually merely wastes of time -- sort of the like the problem of how many angels can dance on the head of a pin, or the question of whether or not Aristotle would be a famous philosopher were he alive today. These well-intentioned but very misguided early

twentieth philosophers, called logical positivists, set out to “solve” all philosophical problems by dispelling them, by re-labeling them as pseudoproblems. Real problems, they said, were scientific in nature; all others, including all classical philosophy problems, were fake problems posing as real problems.

The early twentieth century was a grand and glorious time, even in philosophy -- and we have the logical positivists to thank for it. It appeared for a short time that finally philosophy would make a great leap forward. However, unlike other grand and glorious eras where all things seem possible, the logical positivists managed to make all the profundities seem pointless. The problem of how to treat our fellow humans, for example, was “shown” to be merely the problem of expressing our emotions in the face of what earlier, less enlightened cultures and societies called “moral” and “immoral” acts (this emotivism theory of ethics is sometimes called the “Yea-Boo Theory of Morality”).

Though they meant well, because of such outrageous claims, after a few decades the logical positivists were shown the door, sometimes rather forcefully. The problem of what principles ought to govern how we treat our fellow humans is both a real problem which we have to deal with, as well as, sad to say, a problem that is intractably difficult to solve. In fact, all the classical philosophical problems are once accepted as deeply important and damnably difficult to solve.

Which brings us back to the problem of consciousness. Here is what one eminent philosopher had to say about the problem a few years back:

“The ... most important thing we know about minds is that their states are often conscious. About this, I maintain a gloomy silence.” (Fodor, 1994).

It is the reason for Fodor’s gloomy silence that is the subject of this editorial. I am going to argue that the problem of consciousness cannot be solved because of a logical property of its peculiar nature. Fodor may remain gloomy, but at least he won’t be any more gloomy than he is about the impossibility of perpetual motion machines or the fact that number theory contains true but unprovable theorems.

2. What consciousness is.

The problem of consciousness is the problem of how to explain consciousness in terms of brain states. Consciousness is the way the world seems to us, the way we experience it, feel it. Bite into an onion, look at a rainbow, sniff a dead skunk on a hot summer's day, stub your toe on the foot of the bed frame at four A.M., listen to a baby gurgle and coo. These are experiences, bits of our phenomenology or qualitative feel of the world (philosophers sometimes use the term *qualia* (singular is "quale") to refer to these qualitative feels), and it is experiences, qualia, that give us our subjective point of view. We have experiences because we are conscious. Or rather, our having them *constitutes* our being conscious (you cannot be conscious without being conscious *of* something?) Being conscious is what makes it fun or horrible or merely boring to be a human.

Conscious experience is the most familiar thing in the world. You know nothing as certainly as your own conscious experiences (the contents of your own conscious states), and next to that, nothing as certainly as the fact that you are conscious. This is quite odd given that the inexorable march of science has managed to say nothing illuminating about consciousness. This may seem a bit of an exaggeration, but the statement is true nevertheless. If *science*, not just a few individual scientists here and there, had managed to say something substantial about consciousness, then there would now be some agreement as to what consciousness is and how to explain it. But there is absolutely no agreement on consciousness, neither what it is, nor how to explain it. So, oddly, the thing we as individuals know best, our conscious experience, is the thing about which we as a collective of understanders know least. (I am not saying that science -- in particular, psychology -- has told us nothing about the mind; it has, a great deal, in fact. But it has told us little about consciousness itself. There is a lot more to the mind than consciousness, even if consciousness is its most interesting property.)

The situation is so dire that recently a fascinating book on consciousness appeared which argues that in fact conscious experience is *not* part of the physical world, i.e., that *dualism* is true. I speak of course of *The Conscious Mind*, by David Chalmers (1996), a book that has created quite a stir and that everyone interested in this topic should read. Here it is the late twentieth century, the logical positivists are long gone, and their dream of turning every philosophically deep problem into some sort of physical or biological puzzle has been so robustly overthrown that the universe is now argued to contain not only physical stuff (matter and energy), but nonphysical stuff as well (hence the term "dualism"). Logical positivists hated all nonphysical stuff. What a strange century.

3. The Chalmerian argument for dualism.

Dualism is the thesis that consciousness is not a property of the physical world, that consciousness is not reducible to physical properties of brains or brain states. Chalmers argues that dualism is the correct view of consciousness and its place in the world. The thesis opposite of dualism is called “materialism”. It is the view that consciousness is in fact a physical property. No one currently has any idea which physical property consciousness could be, nor even how it could be a physical property at all. Nevertheless, materialism is the favored view today mainly because of faith in our other sciences.

There is an intuition most of us have, dualist and materialist alike, which Chalmers bases his entire book on. In this section, I am going to show you how Chalmers uses this intuition to argue for dualism. Then, in the next section, I am going to offer an alternative explanation for why we have this intuition that is based solely on a logical property of consciousness itself. My explanation for the intuition shows that this intuition is compatible with materialism being true. My explanation, therefore, renders the intuition useless for Chalmers’ argument, since if the intuition is compatible with materialism, it can’t be used to argue for dualism. In short, I will show why consciousness seems to be *not* a part of the physical world, even though it might well be.

Before turning to the intuition and Chalmers’ arguments, it is important to mention that Chalmers approaches dualism and the study of consciousness in general from a completely nonspiritual, nonreligious perspective. Though he doesn’t play this up, it is one of the most interesting aspects of the book. Here is a dualist who does not in any way fear the mechanistic forces of darkness. In a sense, he wants to be a “mechanistic” dualist; the dualism he advocates is governed by “natural” laws, it is just that these natural laws extend beyond the laws of the physical sciences (and beyond any laws we currently understand). Indeed, he calls his view of consciousness *naturalistic dualism*. Chalmers is an ardent supporter of artificial intelligence and the computational theory of mind. He might even be unique among dualists in believing that the psychological aspects of mind (but not the phenomenological aspects) can be completely explained within the computational paradigm. In short, Chalmers is a dualist who is not a mysterian.

Now, the intuition I speak of is that our conscious experiences could be just what they are regardless of how the world is; that somehow our consciousness needn’t be a part of the physical world. At one time or another, we have all dreamt that we are somewhere strange or vividly

imagined that we are doing something, yet we are not where we dreamed nor doing what we imagined. If we can dream we are eating ice cream in Paris when we are at home in our beds, then perhaps boarding the Concorde and flying to Paris and eating ice cream is also some sort of dream. Extrapolating from this, we all have at one time or another thought something like “What if none of my experiences are real? What if my entire experience of the world is one big dream? What if I am brain in vat somewhere, bathed in wet, warm nutrients merely dreaming that I am reading this editorial? What if nothing is the way it appears to me; what if nothing is the way I experience it?” It is a scary thought. But it could be true. And the fact that it could be true is what Chalmers bases his entire case for dualism on. I call this intuition, this intuition that our experiences ultimately need not cohere with the world we live in, our *Cartesian Intuition* (after Rene Descartes, the great French philosopher and mathematician of the seventeenth century who gave us our first modern formulation of dualism).

Chalmers needs more than just the Cartesian intuition to develop an argument for dualism, for how can the mere fact that we can imagine something mean that dualism is true and materialism false? We can imagine lots of things, and that doesn't mean they are true. We can, for example, imagine a bush on fire yet not burning up. We can imagine that World War Two never happened. We can imagine an unending supply of oil. But burning bushes are consumed by fire, the war did happen, and we will run out of oil by 2050. What Chalmers needs at this point is to refine our Cartesian intuitions a bit. So, from here on out, we need to be philosophically more precise.

The materialist says that all conscious states are brain states, and that whenever those brain states occur consciousness must also occur, just as whenever we have $2 + 2$, we also have 4. The dualist says that though brain states might be the substrate of consciousness, those brain states could exist without producing consciousness. The dualist says that whatever brain states the materialist offers as identical to consciousness, consciousness is in fact a further fact to be explained over and beyond the mere existence of those brain states. According to the dualist, the relationship between consciousness and brain states is not like the relationship between $2 + 2$ and 4, but rather like the relationship between a tomato left out on a kitchen counter and some mold growing on it. There is no necessary connection between the tomato and the mold. The growth of the mold on the tomato cannot be explained merely by referring to the tomato. You also need to mention the mold spores in the air that began growing on the tomato. Or, put it the other way around, you can imagine the tomato sitting out on your kitchen counter with no mold growing on it. So, there is more to mold than merely tomato.

Strictly speaking then, it is not just that we can imagine our consciousness existing independently of or separately from the world, but that we can imagine all the physical universe being just the way that it is but without any consciousness in it at all. Imagining things this way, shows that the physical universe doesn't cause or produce consciousness in any necessary way. Consciousness is sort of *accidentally* tied to the physical universe. This claim requires a bit of philosophical acumen, so I will develop it more slowly and in more detail.

Imagine a glass filled with hot water. The atoms in the glass are caroming all over the place in a very agitated way (let's call the behavior of these atoms "low level facts"). Now, try to imagine another glass of water where the atoms are behaving in *exactly* the same way as in the first glass, but where the water in the second glass is cold. You can't do it. Or, if you think you can, you are mistaken. For, all we *mean* by "hot" is that the atoms are caroming all over the place in a very agitated way. Fix the behavior of the water atoms in the glass and you automatically fix the water's temperature. This example exhibits just what is going on in our entire physical universe. It is simply inconceivable that the low-level facts about our physical universe (the atomic and subatomic facts) could be what they are and yet there be no stardust, no suns, no galaxies, no planets, no continents, no minerals, no life, no United Nations, no penguins in Antarctica, and no MTV (the Music Television Channel). In short, and though it may sound strange, MTV is what it is merely because certain low-level physical facts of our universe are what they are. There is no imaginable universe with the same atomic facts as ours that isn't blessed with MTV.

Another good example is your computer. Fix the states (1 or 0) of all the bits in your computer (in the CPU, on all the RAM, and on your hard drive), and you thereby establish the program you are running (a word processor, say), what you are doing with it, and what things are stored on your hard drive. There is nothing more to the information in your computer than the states of the bits that make up your CPU, RAM, and your hard drive. This editorial which I am now writing is not something over and above the bits that make it up. Change even one of those bits and you change the editorial (probably by introducing a misspelling). The bits are the low level facts, and the editorial as it appears on the screen as I write it is the higher level fact.

Our physical universe as a whole works just like the computer (or the glass of water). Since fixing all the low level facts in our physical universe completely determines what all the high level facts are (note, this doesn't go the other way), in a sense our universe is just one big set of physical facts described at different levels: the lowest level is the subatomic level (or the sub-subatomic level), from there we move up through the atomic level, the molecular level, the

macrochemical level, the biological level, and on up to the level of cultures and politics. This is a beautiful picture, and the one supplied to us by our modern science, but notice that *consciousness needn't be a part of this picture*, for we can imagine fixing all the low level facts of our physical universe without thereby introducing or establishing consciousness at all.

Here's one way to do this imagining: imagine a physical universe just like ours but populated with perky zombies (not the Hollywood version) or lively, bouncy somnambulists who do just what you do as they go about their day, but they don't actually experience or feel anything. They elect governments, smile at penguins, and watch MTV, but they are completely inert on the inside; they feel neither a sense of community, nor amusement, nor bemusement. Though it might seem hard to imagine this, with some care, it is *coherently* imaginable, which makes it radically different from the case of imagining two glasses containing identically behaving molecules of water but the water is at different temperatures. This last imagining is *not* coherently imaginable (imagination is not limited even by the laws of logic, which is both a blessing and a curse).

The crucial part of Chalmers' argument is right here: we can coherently and logically imagine a physical universe exactly like ours peopled with lively, bouncy, but completely unconscious somnambulists, but we cannot coherently and logically imagine a physical universe exactly like ours without electrons, atoms, molecules, insects, penguins, and governments (just as I cannot coherently imagine a computer exactly like mine here with exactly the same states in all of its bits but not running MacWrite and not storing this editorial). Because of this difference in what we can coherently and logically imagine (these two adjectives are crucial for philosophical precision), it must be that consciousness is different from electrons, atoms, molecules, etc. In fact, Chalmers concludes, it must be that consciousness is not really physical after all. Consciousness, he says, seems to be a further fact over and above the physical facts of the universe.

4. Why dualism seems true, even if it isn't.

In his book, Chalmers gives several versions of the above argument. The Cartesian intuition funds them all because imagining a physical universe identical to ours without consciousness in it is possible solely because our consciousness feels separate from the universe, i.e., that our consciousness could be just what it is even if the universe were very different from

what it is.

Some philosophers and cognitive scientists are persuaded by Chalmers' arguments and some aren't. Some deny that they have the Cartesian intuition (how they do this is a mystery to me -- I suspect they are in denial), and some agree that they do have it but deny its importance. No one can gain the upper hand. Which is just saying that the problem of consciousness is as deep and perplexing as it ever was. Even on an unsympathetic reading of Chalmers' book, he has demonstrated for us once again just how deep and perplexing the problem really is. But suppose he is right?

That dualism might be true is profoundly troubling. Dualism, especially Chalmers' version, has a certain beauty to it, but it is completely at odds with our late-twentieth century scientific, materialist worldview. That something so important to humankind -- the fact that we are conscious -- might not conform at all with the other facts of the universe has to strike anyone as alarming (after years of being thus alarmed, one tends to get gloomy, hence Fodor's mood). But it seems to me that the best strategy for arguing against dualism is not to try to argue directly that it is false (because, as I said in section 3, we all have the Cartesian intuition that our conscious experience could be just what it is independently of the world), but rather to explain why we have the Cartesian intuition *even if materialism were shown to be the correct theory*. This is what I will now try to do.

Here is my strategy. We all agree that we do in fact have the Cartesian intuition. I will show that conscious experience itself leads to this intuition. If this works, then it will forcefully undermine our confidence in the Cartesian intuition as an argument for dualism. Why? Because if consciousness itself leads to the Cartesian intuition, then the intuition that consciousness is logically separate from the physical facts of our universe is due not to its being in fact thus separate, but due solely to the fact that we are conscious. In other words, materialism could be true, but our consciousness would be exactly as we experience it now, in particular, we would still experience consciousness's peculiar property of seeming somehow independent of the universe. So, I am claiming that one of the essential properties of consciousness is the feeling of separableness (and separateness) it engenders in us (the Cartesian intuition). This feeling or intuition exists in us provided only that we are conscious. It follows from this that the feeling or intuition cannot be used to fund arguments about universes without consciousness, i.e., it cannot be used in arguments for dualism.

Since on the materialist assumption the Cartesian intuition should not exist, the best way

for me to make my argument is to assume that materialism is true and show that we would still have the Cartesian intuition, i.e., that our consciousness would still feel separable from the universe we inhabit. So, first assume materialism is true, i.e., that the relationship between your consciousness and your brain processes which give rise to it is not like the mold growing on the tomato, but rather like the relationship between the states of all the bits in a computer and the program it is running. As I discussed in section 3, this relationship is so logically tight that in fact there is just one thing merely described in different ways, exactly like $2 + 2$ and 4 . To help with my exposition, I introduce a special term for this assumed materialistic relationship, between brain processes and consciousness, of being logically identical but descriptively different. I will call it *double-aspect identity* to stress that there is just one thing describable in two different ways (or aspects), both as a brain process and, at a different level, as consciousness.

Here's my argument. What would you experience if you saw the brain process which are double-aspect identical to your consciousness? You would just experience another quale. That is, you would have just another conscious experience. You would not see that brain process cause or result in your conscious experience because, as just pointed out, you are already conscious of perceiving your brain process. It is impossible that you could experience something causing or resulting in your experience, that you could be conscious of something causing or resulting in your consciousness, because when you turn your attention to that thing, you are already conscious of it. Put another way, you cannot go from being unconscious to being conscious and at the same time perceive the processes which result in the transition. Why? Because by the time you perceive the processes that result in the transition you are already consciousness. Just as you cannot hear the clapping of your hands before you clap them, you cannot be conscious of what's causing your consciousness before you are conscious.

But you can go from being asleep (unconscious) to being awake (conscious) and determine that an alarm woke you. Isn't that seeing (or perceiving) what caused your consciousness? No. Your consciousness does not arise from the workings of an alarm clock. It's true it woke you up, but your consciousness is not double-aspect identical to the workings of the clock. Even on a dualist assumption it is brain processes and not anything else which results in consciousness (albeit, not in any logically tight way).

Here is a more detailed version of my argument. Call the physical, neurological processes to which our consciousness is double-aspect identical the "consciousness producing processes," or CPP for short. The relationship between CPP and consciousness is exactly like $2 + 2$ and 4 .

Now,

what would you see (what visual experience would you have) if you saw your CPP? You would see some working neurons; you would experience just another quale. That is not seeing or experiencing your consciousness arising from or depending on your CPP. It couldn't be. Why? Because experiencing your CPP is not experiencing it as giving rise to your consciousness; it is merely having yet another conscious experience. Experiencing your CPP presupposes that you are antecedently conscious, hence your CPP has already done its job.

Consider this question: How could we see your CPP result in our consciousness? In order for us to see CPP actually result in our consciousness, we would have to see CPP causing our consciousness. But we can't do this because seeing CPP is having a conscious experience in the first place (of CPP -- a brain process). We are, as it were, forever locked inside our consciousness. We can't see anything causing it or resulting in it because we can't step outside our consciousness and watch it come into existence via some causal process. But being locked forever inside our consciousness is what is responsible for the Cartesian intuition in the first place. It is this being locked inside our own consciousness which gives us the intuition that the world, while relevant in fact to our conscious experience, doesn't seem necessary to our conscious experience.

And this just is the Cartesian intuition. Hence, even if materialism is true, we would still have our Cartesian intuitions. And hence we would still have the intuition that Chalmers uses for his argument for dualism. And hence dualism will still *seem* true. So the intuition cannot be successfully used to argue for dualism because the intuition is compatible with the opposite of dualism, namely materialism. So, Chalmers arguments are stymied.

Notice that I have not given a positive argument for thinking that materialism is true, nor even an argument for thinking that dualism is false. I have merely shown that one good, modern argument for dualism, Chalmers' argument, won't work because its basic assumption is compatible with both dualism and materialism.

What does all this mean for us AI types? Consciousness remains as mysterious as ever. We don't know whether materialism is true, nor whether dualism is. We don't know what brain processes result in or cause consciousness. We don't even know if other things are conscious: insects, birds, cats, trees, even other humans (you seem conscious to me, but maybe you are just sleep walking). I think we are entitled, with Fodor, to be rather gloomy about this state of affairs. The most profound thing about being intelligent and cognitive is also the deepest, most perplexing puzzle we have. Will we ever know what consciousness is and where it might be

found? I don't know. But I for one am going to think twice about turning my computer off.