

Debate on Neuroscience and the soul on <http://forums.philosophyforums.com>. Quotations are in red and the responses by Faustus (Brian Petersen) and Death Monkey (Kevin Dolan) are in black. Note that sometimes a quote (in red) contains a previous response from Faustus or Death Monkey. This is usually found at the beginning of the quoted portion (in red) and separated from the actual quote by a horizontal line. This is usually done to provide a context for the quote before responding to it.

Here is a summary of the model that appeared in a special issue of the journal *Cognition*, which was dedicated to examining the model's various features and challenges:

At any given time, many modular cerebral networks are active in parallel and process information in an unconscious manner. An information becomes conscious, however, if the neural population that represents it is mobilized by top-down attentional amplification into a brain-scale state of coherent activity that involves many neurons distributed throughout the brain. The long distance connectivity of these "workplace neurons" can, when they are active for a minimal duration, make the information available to a variety of processes including perceptual categorization, long-term memorization, evaluation, and intentional action. We postulate that this global availability of information through the workplace is what we subjectively experience as a conscious state.

In other words, your brain is a Swiss Army knife affair with various modules dedicated to processing very specific types of information. These modules are networks of neurons which are sensitive to certain modalities of information content, and tend to respond only in the presence of that sort of content. Take vision, for example. Early on in the processing stream, there are modules that do nothing more than detect edges, luminescence, color, and other very primitive, raw aspects of the input from the eyes. Further on, different modules take the "judgments" made on at the lower levels and process the information at a much deeper level of sophistication—recognizing shapes as letters, collections of letters as words, shapes as being faces, faces as belong to specific people, and so on.

As these various modules come to their conclusions, they broadcast their results (via neuronal signaling, of course) in a kind of temporary memory nicknamed the "workspace". These broadcasts can differ in strength, and on occasion conflicting or contradictory conclusions will be reached. At any rate, the content in the workspace can then be processed and reacted to by other modules and decision making areas—so long as that content is in the workspace in sufficient strength and for a sufficient length of time. How strong and how long varies according to the type of information being processed and various other factors.

When information content is in the workspace and worked on by a sufficient number of areas—especially those associated with memory and motor functions—then subjects will report awareness of the content. In the case of conflicting or contradictory judgments, many sensory modalities operate according to the principle of "winner takes all"—only one conclusion at a time will be reportable by a subject, even if parts of the brain come to a different judgment.

There is no finish line in the brain demarking an instant or objective point when content becomes conscious. Awareness of content is a matter of degrees. At one end of the spectrum, subjects will report no awareness, in the middle, they actually may not be sure but can guess at better than chance, and at the extreme other end, we have full, robust consciousness.

By Faustus (Brian Petersen)

Quote:

Originally Posted by **Tagfat**

Sadly it is not unusual for authors of neuropsychological pop-science articles to grossly overstate the case that can be made from the available evidence.

Authors of neuropsychological pop-science articles might do this when they are merely journalists, but in the case of Pinker, we are dealing with a cognitive scientist, someone in a very good position to understand what the available

evidence actually warrants. Rather than being “pseudo-science”—a ridiculous charge—the ideas he expresses in the article are bland and quite well accepted in the mainstream of scientific thought, where the idea that people have souls has been rejected for over a century by an overwhelming majority of those familiar with the evidence.

Quote:

Originally Posted by **tagfat**

There isn't really anything new in linking states of mind to physiological states, that has been done since the days of Descartes. . .

This is a preposterous claim when you pay attention to what Pinker was actually talking about--he was discussing techniques of brain imaging that have only really begun to take off in the last decade or so. With these, we have been finally able see what parts of the brain are responsible for processing very specific types of information, and this has enabled neuroscience to have successes that would have been unthinkable back in the early 1980's, let alone Descartes' time.

Quote:

Originally Posted by **tagfat**

. . . but it has long been known that these correlations are not as simple as eliminative materialist "psychologists" - like Paul and Patricia Churchland - would have you believe.

First, they are not mere “psychologists” but work in the field of neurology, emphasizing (particularly in Patricia’s case) the philosophical implications of their findings. They are experts in neural network and connectionist models. I don’t always agree with them—I think eliminative materialism throws the baby out with the bathwater--but they are must-reads if you want to be educated in this subject.

Second, where have either one of them ever put forward the notion that correlations between mental states and brain states are “simple”? This does not jive with anything I ever remember reading in their books or papers and actually seems to go against the main thrust of their work. Can you provide a reference that would constitute evidence that the position you ascribe to them is actually what they believe?

By Faustus (Brian Petersen)

Quote:

Originally Posted by **Monroe**

Then I don't see why you used it in a reply to my philosophical concern that "The fact remains that neuroscience doesn't explain how the brain produces consciousness, nor is there any theoretical explanation currently available (nor conceivable for many)."

That’s because I mistakenly believed you were making a claim about science. Taken that way, your statement is just false.

So you have a “philosophical concern” about a supposed “fact”—that neuroscience doesn’t explain consciousness and there aren’t any theoretical explanations available.

First off, not all philosophers would agree that this declaration of yours is a “fact” at all. Facts tend to be the sorts of things about which there is widespread consensus. There is no widespread consensus in favor of the mysterian position. It is merely an opinion, as is its denial.

However, there is a consensus in the scientific community that the workspace model broadly understood is the correct way to explain consciousness. Indeed, this agreement is shallow in places, since there are all kinds of disputes over the details, and different researchers tend to emphasize different aspects of it. But the thumbnail sketch I provided is solidly acknowledged and highly unlikely to be abandoned by any sort of revolution.

So here we have two parallel streams—scientists doing their thing, and philosophers doing theirs. If the explanations offered by scientists are not satisfactory to some philosophers, what am I to make of this? Should I be concerned? If

a theory succeeds as science but is doubted on purely philosophical grounds by some (not all) philosophers, has the theory failed? I don't think so. If merely philosophical concerns about explanations for consciousness cannot be expressed in a manner that has scientific relevance, then those concerns are a sideline and I will dismiss them as idle academic nonsense.

Quote:

Originally Posted by **Monroe**

A philosophical worry is, what exactly is information content (probably the easier problem), and how can a physical object have it, in a way that is intrinsic and not by fiat by a homunculus (like when we bestow information content to written language by fiat)? So I don't think your dismissal of philosophy will work, unless you are willing to just ignore issues with the basic assumptions of your viewpoint.

Remember, we aren't talking about a viewpoint that is unique to me. We're talking about the informed opinions of the scientific community, which I am defending.

So your philosophical worry is how physical objects can have information content in a way that is intrinsic. The first part of your concern (what's information content?) isn't merely a philosophical concern, it is a legitimately scientific concern. The second part (the requirement that it be "intrinsic") is not. The demand for intrinsic meaning strikes me as the sort of purely philosophical point which scientists can justifiably roll their eyes at and ignore. At least, I'll be inclined to dismiss it until I hear an independent argument demonstrating a) why information content **MUST** have this property, and b) how such a property can be consistent with scientific models of biology.

Any legitimate scientific theory involving a concept of informational or representational content in a biological system is going to have to face up to the consequences of evolution by natural selection. This places two constraints on us:

1) Whatever information content is, the mechanisms responsible for it must necessarily have had their origins in systems that had no such content.

2) Representational capacity must have evolved bit by bit, so that at one end, we have non-living objects like rocks with no powers of representation, at the other end we have human beings with language. The ability to represent must then have happened one tiny step at a time, each step conferring some sort of benefit to the organisms possessing it, with the first few steps barely even qualifying as representations, but falling somewhere on the edge of the concept. This is how all complex functions evolved.

I cannot imagine how the demand for intrinsic meaning can be met in light of these constraints, but that could just be a failure of my imagination. At any rate, when philosophical constraints butt heads with scientific realities, I side with science every time.

Here is an example of information or representational content as it exists in a very primitive form: quorum signaling among bacteria. Some bacteria secrete non-essential compounds and have receptors on their cell walls which can detect the level of the compound in the immediate environment. When the levels reach a certain density—which in normal conditions "means" that the population of bacteria is at a certain density—a process of gene transcription in the bacteria will be activated, and proteins synthesized. The end result is a group coordinated activity beneficial to the population of bacteria. (The activity depends on the organism, but in each case quorum signaling results in coordinated group behavior.)

Thus, quorum signaling is an adaptation which enables bacteria to detect and represent information about population levels so that this information can be used to trigger beneficial behaviors. It is the stuff of which all capacities to represent are made out of.

Quote:

Originally Posted by **Analytic**

I never understood what is the exact content of the claim 'the brain is a computer'.

It merely means that the brain performs computations on inputs. Instead of using bits turning on and off as in a PC, it uses connection strengths between neurons in vast networks. Many neuroscientists are rightfully hostile to the brain-as-computer metaphor, because on the surface at least, it appears to have been taken too literally by some. It

isn't clear, for instance, how the software/hardware distinction can be maintained given what we know about the brain. But the brain-as-information processor model is well entrenched in various forms, and that's all Pinker is talking about when he discusses computation by the brain.

Quote:

Originally Posted by **Analytic**

The questions prompted by Searle's Chinese Room thought experiment are still unanswered.

You write as if Searle has discovered some sort of fundamental law or truth, and that researchers in cognitive science need to stop what they are doing until they have satisfied him and his followers. Are you not aware that virtually no one involved with the science of consciousness was ever convinced that his argument has merit? Most of us think his arguments have already been answered, and were demolished in the very issue of *Behavioral and Brain Sciences* where which they were first made. We've moved on.

Quote:

Originally Posted by **tagfat**

What makes it pop-science is the reference to impressive technology (impressive to the layman) while the central questions of the real debate is omitted.

No, what makes the article pop-science is that it appeared in a news magazine and did not contain new ideas or standard scholarly references. Whatever these mysterious "central questions" that were omitted are, I'm sure they have been addressed in the appropriate forums, which would not include *Newsweek*.

Quote:

Originally Posted by **tagfat**

What probably makes it pseudo science is the heavy financial interest in the view and the allied view that "depression is a serotonin imbalance in the brain", a cornerstone in the marketing of anti-depressant drugs.

Absurd. Pseudo-science is not defined by how well financed a view is, but how well supported by scientific evidence it is, and how closely a researcher sticks to the evidence. The ideas in the article are not controversial among scientists, and would probably only be challenged in Internet debate forums rather than journals dedicated to science.

Quote:

Originally Posted by **tagfat**

If you think that you have evidence that nothing like a soul is involved in consciousness then you should present it.

You've got the reasoning requirements backwards. The burden of proof is entirely on dualists to offer evidence to support the existence of something they believe in. Since they have had absolutely no success in this, their beliefs are now essentially dead in the water.

Re the Churchlands: you have not produced any evidence that they claim correlations between mental states and brain states are "simple," which was your original claim.

Re the supposed amazing paradox of van der Waal interactions: sounds like you've been suckered by Stuart Hammeroff, whose ideas on this subject—by his own admission—are not well regarded by his peers. Even if he's right about quantum effects in microtubules (highly unlikely) this would merely be a variation within the current paradigms of brain science, and would lend no support whatsoever to dualism. The truth is that there is no paradox involved. Brain activities aren't merely "chemical"—they are electrochemical, and it's no surprise that effects which disrupt or change patterns of electrical activity would have an effect on consciousness, even if RIGHT NOW, we can't model why a particular gas has the effects it has.

By Faustus (Brian Petersen)

Monroe,

Quote:

This doesn't answer "how." How are the properties of consciousness physical properties?

In the case of reducing water to H₂O, it's not necessary that we know all the properties of H₂O molecules and how their group behavior can be explained by those properties before we make the identity claim and the claim that the properties of water can be explained eventually by this molecular theory. This is because we already know that the macroscopic properties of water that we are trying to explain are of the same fundamental kind as the microscopic ones. We are just trying to explain the whole by the parts. The molecular properties of electrostatic interaction are simply about how masses relate to each other in a space-time structure, and this is what we want to explain on the macroscopic level, just on a bigger scale.

The key point here is that they are the same "type" of properties, namely physical properties. But what does that actually mean? If consciousness affects other physical things, then how is it any less physical than anything else we call "physical"? What other requirement do you think being physical has?

Quote:

But it is not clear that the "macroscopic" properties of consciousness (I mean the ones we observe by being conscious) are just physical.

Sure it is. Every single one of them has observable effects on the physical world. Therefore they are physical. What is not necessarily clear is that they are reducible to the laws of physics as we currently understand them. That is a completely separate issue, though. Consciousness is clearly a physical phenomena. The only thing that is really in question is the mechanism by which it works.

Quote:

Surely they are not just the macroscopic physical properties of the brain: I can look at a brain and not be able to tell you anything about its consciousness.

So what? I can look at CPU, and not be able to tell you anything about the computational processes going on inside. All that means is that (1) I don't understand how it works, and/or (2) I lack the necessary equipment to extract the information that is there. Neuroscientists most definitely *can* look at a brain (with the right equipment) and tell you something about its consciousness. Not everything, but enough to render your above argument completely null and void.

Quote:

So obviously consciousness is not a macroscopic physical thing, because all macroscopic physical things are third-person observable by humans; that's part of what it is to be macroscopic. So how then are we to explain consciousness via physics? It is not reasonable to just assume that we can since it's not clearly that consciousness is even in the categories of physics.

On the contrary, every aspect of consciousness that you could possibly know you have, has some macroscopic observable effect on your behavior. There are exactly two options. (1) Consciousness can, at least in principle, be explained in terms of physics, or (2) It cannot, in which case those physical effects cannot either, ie consciousness is supernatural. There is no evidence to support the latter option. In any event, claiming that consciousness is not physical is not just an unsupported claim. It is demonstrably wrong.

Quote:

If you want an alternative theory, I think maybe there is more out there in the world than the categories of physics cover. Perhaps mentality is an intrinsic feature of the world, observable only "from the inside" so to speak. Like "quantum consciousness," I think.

The above is simply not coherent. First of all, there is the physical effects issue I already mentioned. Whatever the mechanism of consciousness is, it qualifies as a physical phenomena every bit as much as any other feature of the world which we call physical. Second, quantum mechanics is physics. The various proposed ideas of how quantum mechanics might play a role in consciousness, all place consciousness quite firmly in the category of "physics".

Quote:

And metacristi, this is not a problem of formulating theories that make better predictions, or formulating ones that are best supported by the given laboratory evidence, or figuring out which current conjecture looks most probable in light of empirical evidence, or of the materialists proving definitively that their theory is true. No, the problem is conceptual. It is about, first of all, trying to come up with a theory (not even necessarily true or supported by any experimental evidence at all) that could explain consciousness via physics. It is a theoretical problem, a problem solvable only by armchair methods. Once we are able do that, then we'll have an idea of how we are to proceed in finding a correct theory via scientific method.

The ironic thing is that while armchair philosophers are wasting their time trying to figure out whether it is possible in principle, actual scientists are *already doing it*. Neuroscientists already understand quite a bit about how many of the various phenomena which we group under the vague heading of "consciousness" work. The fact that every time our knowledge increases, the dualists simply move whatever aspect of consciousness has been understood over from the "consciousness" category to the "neural correlates" category, just demonstrates that they are not really interested in understanding anything, or solving any problems. They are interested in protecting their irrational belief that there is something more to their minds than what a lump of meat in their head is doing.

Nothing short of a complete reductive explanation of how every aspect of the mind arises from neural activity, is going to satisfy the dualists. And even then, you will still have epiphenomenalists claiming that something is still being left out, even though their own philosophy implies that whatever it is, they could not possibly know about it.
By Death Monkey (Kevin Dolan)

Quote:

Originally Posted by ~E

Isn't it a fairly well established result that an observer is required to determine events? Without an observer aboard, which seems to entail consciousness, zombies can't determine anything, ne?

I am aware of no such established "result," but then I'm actually unsure of what you meant to say. Science takes all the properties and activities of the whole observer and breaks them down into processes happening at different times and places in the brain. These in turn would not have properties we'd associate with observers or conscious states. If they did, we'd be using circular reasoning instead of actually explaining anything.

To Analytic and Muxol regarding my little stumble over the meaning of "computation": what I was trying to do was pay respect to the fact that the term is often used very loosely. I didn't want to commit myself to some narrow definition or other. When someone talks about information processing on the surface of a black hole, the concept has a relationship to information processing in a brain, but they aren't the same claim. When someone talks about computation by neural networks, many of them at least will try to convince you that this has zip to do with Turing Machines. So I was trying to be as neutral on all those issues as humanly possible.

Quote:

Originally Posted by Analytic

Your posts strongly suggest that the question of consciousness has already been solved by empirical, scientific methods and it has a theory that is not speculative any more (minor details excepted), so it's like the standard theories of physics, say. Am I right in this? Has neuroscience really attained the standards of physics?

I did not mean to make so strong a suggestion, and I certainly don't think anyone should hope or try to make neuroscience like physics. The techniques and tools of physics do not now and never will resemble those used in biology, and physics envy is to be avoided. This is not a problem with standards for biology as a science, just a truth about how different tools are used in different domains of knowledge.

The biological realm has irreducible issues of interpretation with regards to what various functions are "supposed" to do. It involves reverse engineering in ways physics does not. Mass is can be measured directly. Whether the fins of a stegosaurus evolved for cooling or other purposes cannot. (The only way to confirm or disconfirm such theories is to find out how well a structure would work at a proposed function and model how likely it is to have evolved for that function.) These issues extend right up to the brain and become entangled with issues of intentionality and consciousness. They have to be incorporated and cannot be made to go away by making a "better" theory that somehow resembles physics. This is why reductions from mental state terminology to precisely defined physical states are impossible.

So in answer to your physics-envy style questions:

1. Is there an agreed upon unambiguous operational definition of consciousness as such?

Nope! Don't expect one, ever. Our mentalistic idioms are more like "game" than they are like "mass". We can't make a definition of the concept "game" which is immune to borderline cases, either. But that hardly prevents anyone from playing and designing games. Everyone can agree that if a subject correctly reports "I have just seen the word 'tiger' flashed in my right eye" that an unambiguous conscious state occurs. It's the weirder cases you get with brain damage or where a subject reports nothing but nonetheless parts her brain recognized "tiger" on their own that we don't know what to say. Solving these border cases is about deciding what sorts of concepts we want to have rather than inventing a better theory that will capture some hidden truth in nature.

2. Is your definition of consciousness all-or-nothing, or does it allow for quantized values?

I would say it is neither all or nothing nor able to be quantized, if by "quantized" this means that consciousness comes in units that an instrument could ever conceivably measure. (Again, as we slide from cases of activities clearly not a game to those that clearly are, it isn't as if we are measuring some sort of "gameness" value that adds up until a language-independent threshold is reached.)

3. Do you have testable theories concerning consciousness? In other words, do you have a well-defined method by which you can classify physical systems as those having consciousness and those having none?

None are needed, since natural selection has designed most animals with sophisticated machinery able to accurately recognize and model the mental states of other organisms and use those models for prediction and deception. Even ravens have this capacity, as recent experiments in animal cognitive ethology reveal.

What this indicates is that the truth conditions for the possession of mental states are made independently of anything science tells us about what goes on inside of brains. Except in very bizarre circumstances, asking for a scientific test for consciousness is exactly like thinking that you validate a checkmate by a chemical analysis of the pieces and board. In normal circumstances, it is the patterns of macroscopic behavior which confirm or contradict attributions of particular states, and no discovery about brain events could trump evidence coming from that level. Science can only address the causal mechanisms that lead to those patterns, so a testable scientific theory will be something along the lines of "Spiked activities in brain areas X, Y, and Z will be present in all cases where subjects report seeing the word 'tiger' flashed for 150 milliseconds in one eye". Or something like that. And yes, there have been successes along those lines.

However, I don't know of any cases where ALL the necessary and sufficient brain process for specific kinds of reports have been isolated. Given the context dependence of brain activities and their chaotically complex nature, this might be impossible for perfectly banal reasons having nothing to do with metaphysics. That's where the challenges are.

4. Do your theories have predictive power in the sense that you can predict of a physical system through inspecting its structure whether it has consciousness or not?

This seems like a rewording of question 3, so I'm sure you can guess the answer: no, because the structure isn't the place to look in normal circumstances, though the patterns of behavior resulting from that structure are.

In highly unusual contexts, for example a woman under complete physical paralysis, looking at the structure might be appropriate. We could right now measure the extent to which her brain is processing the effects of hearing words or sounds. If she can hear "Raise your right finger for 'yes'" but cannot act on that request because the relevant damage is in the motor areas, we might be able to watch outputs to them in response to questions. If those outputs and activities in brain structures made sense as responses to the questions put to her in just the way watching her move her hand would, then we could perhaps conclude she was conscious. In this case, all we are doing is substituting a suitably interpreted brain event for behaviors we normally see elsewhere in her body.

The only problem with this scenario is that we can see outputs to motor areas in cases where we've fooled the subjects into not being aware of a stimulus, so we know that the brain can process not only the semantic meaning of a stimulus, but also respond to a prior motor request conditioned on that meaning, yet do so without the subject

being conscious of it. You have to create some bizarre contexts before this happens (brief flashes of stimulus followed by a masking stimulus), but thought experiments about brains in vats and weird paralysis are just these kinds of contexts. So right now, we don't know enough to be able to carry out the experiment I described, though it isn't outside the realm of possibility and we're close to pulling it off.

5. Lastly, do your theories make it possible to create physical systems in laboratory that pass your tests of consciousness? If so, what is their complexity?

Yes and no.

Yes, in the sense that attributing intentional states to artifacts and animals is justifiable to the extent that it helps to accurately predict and model their behavior, and we already have AI rich enough in its behavioral capacities that people naturally find themselves making such attributions (“the chess program believes I’m about to checkmate,” “don’t get too close to the virtual soldier or he’ll see you”).

No, in the sense that even the best AI justifies only a very, very barren version of such attributions. Unfortunately for fans of the Turing Test, people tend to err on the side of attributing far more depth and complexity to animals and AI than would be justified by a truly careful examination. Simple text based AI have fooled some naïve people into thinking the programs are much smarter than they are. This demonstrates that the Turing Test has limitations, and really could only work with a very well informed and trained panel of judges, since arguably some AI has already passed the test with folks who don't know what sorts of questions would trip the AI up and reveal its limitations.

No also in that the AI used to generate these minor successes is not related to neuroscience in most cases.

But the Cog Project and the works of Rodney Brooks are pretty interesting places to look for future developments.

All that aside, my larger point about harping on the workspace model is to point out that there is a consensus about it as a valid scientific model of consciousness, even if the consensus exists only at a very abstract, “big tent” level. I’m doing this because of the blatant exaggeration one continually encounters on the part of mysterians who keep trying to insist that science has made no progress on consciousness, which is just baloney.

The bottom line once again is that Pinker was justified in ALL the assertions the article made, which were utterly bland.

I should point out in all fairness that he is far more skeptical than I about the successes we’ve already achieved, since he tends to side with people like Chalmers about the so-called “hard problem”, which I will continue to insist is a philosophical phantom. But you’d never know that from the article alone.

Quote:

Originally Posted by **Monroe**

When he gets the stimulus producing the bear experience, he believes there is a real 3-dimensional bear, existing in a physical 3-D space. So this aspect of his belief doesn't have a causal connection to the thing it represents.

Sure it does. The structures in his brain are representing a virtual bear, with whom they have a perfectly causal relationship. He thinks the causal relationship involves a real bear, but he is wrong. This doesn't mean the representational activity happens by magic, uncaused by any event in the real world. The falsity of his beliefs revolve around the nature of the causes.

Quote:

Originally Posted by **Monroe**

To put it another way, you could have the same input being a brain in a vat, or being a normal body, and have the same beliefs and representations.

Agreed. I made the same point when I said that what happens physically between input and output doesn't change whether we are in vats or not.

Note that you move from this to discussing whether the representation is “right” or whether his resulting beliefs are

“false”. Now we have left the realm of science and entered philosophy.

What sort of label we choose to assign to brain activities in this sense has no scientific content, and is merely a sort of conceptual or legislative move, one of frankly limited usefulness or consequence. Perhaps if we found we were brains in vats, we would want to use a different word for activities in the brain that represent virtual events from those which represent real events. Or maybe we'd start tossing around concepts like “virtually true” or “virtually false”. Perhaps we would be motivated to create these distinctions if virtual reality technology became widespread, and maybe there would be legal consequences for these distinctions. But there would be no scientific consequences, nothing to change how neuroscience is done.

Same goes with the issues around what causes Bigfoot-beliefs: your concerns are about what we want to call various kinds of epistemological relationships, and these are not concerns which bear upon brain science.

Quote:

Originally Posted by **Monroe**

On the other hand, if your position is that we can't have beliefs about things other than our perceptions themselves, then I can see how you would still disagree with me. But that position is obviously false.

Agreed, since representations which cause consequent belief states are themselves caused somehow, even if not by the things which a subject may think they are caused by.

Quote:

Originally Posted by **Monroe**

So is consciousness in the same categories as the categories of physics? Are the properties of consciousness a subset of the various kinds of mass and force and spacetime in physical theory? How?

Well, given my long anti-reductionist, anti-physics envy reply to Analytic. . . .

By Faustus (Brian Petersen)

Monroe,

Quote:

DM, I think your position is basically that the mind, as experience first person, is not distinguishable from its behavior or "physical effects." This is just straight up eliminativism.

No, my position is that there is no aspect of your first person experience which does not have physical effects, and which therefore could not be studied scientifically. This is not different than saying that we know of no aspects of electrons which do not have physical effects. It is no more saying that the those experiences are identical to those effects, then it would be to say that an electron is identical to its physical effects.

There is no aspect of your first person experience which does not have physical effects. There is therefore no reason, short of appeals to supernaturalism, why we should not be able to study and describe consciousness in terms of those physical effects, in exactly the same way we do for every other physical phenomena.

Quote:

I can't debate it. Maybe you don't have consciousness. Whatever, we've talked about this before to no avail. You don't seem to grasp the difference between a thing and the effects it produces.

I grasp the distinction just fine. You do not seem to grasp that it is not relevant. You are talking about metaphysics, and I am talking about epistemology. Science describes things in terms of their physical effects. **If you wish to claim that consciousness cannot be described in terms of its physical effects because it is supernatural, then you need to justify that claim. If you are claiming that such a description would be incomplete because it would leave out any non-physical components of consciousness, then you need to explain how you could possibly know that your, or anybody else's consciousness, has any non-physical components. You might as well claim that electrons have a non-physical properties. Maybe they do, but how would we ever know? And if we cannot know, why should we care?**

Quote:

As for your CPU example, to understand the computational aspects of the CPU you'd have to understand the purposes for which it is designed. Anyway, its not relevant to my minor point which was that consciousness is simply not a macroscopic physical thing.

A point which I quite clearly refuted. The very fact that your consciousness affects your behavior in such a way as to make it possible for you to talk about it, refutes that point.

Quote:

Computation by computers is basically a matter of input and output corresponding to some desired rules. If you believe that that's what the mind is a matter of, I disagree.

Disagree all you like. Do you have anything other than your opinion to back up this belief?

Quote:

I experience an inner life that is in between the input and output. I don't experience my brain processes as the neurologist views them, I experience my consciousness.

So what? This has absolutely no relevance to my points. None of this is even remotely incompatible with the hypothesis that your consciousness is something which your brain is doing.

Quote:

The materialist neurologist must show that the subjective aspects of the mind are deducible from the physics-describable aspects of the brain. But how could he do such when they are not physics-describable?

What on earth do you mean by non-physics-describable? Do you mean that they have no physical effects? Because that is simply false. Do you mean that they cannot be described in terms of those physical effects? If so, why not? What, other than your intuitive preconceptions, indicates to you that they cannot be?

What is it that you think makes consciousness so different from other physical phenomena? In what way do you think that the process of scientifically studying and understanding consciousness in terms of its physical effects, will fail? So far it appears to be working quite well.

Nomos,

Quote:

Solipsism is unfalsifiable. It follows that it is not possible to show that consciousness arises from matter.

Solipsism is unfalsifiable, which is exactly why it is utterly vacuous as a position. In any event, this only means that it is impossible to *prove* that consciousness arises from matter, which is also irrelevant, because that is not what scientists try to do. Scientists do not need to prove that consciousness arises from matter. What they need to do is construct a scientific theory which describes how consciousness works. You are talking metaphysics, which science has absolutely no interest in.

It could be that the entire physical world is a construct of minds, and that the brain is just the projection of those individual minds onto that artificial reality. So long as this imaginary physical world functions naturalistically, including those projected imaginary brains, science is perfectly happy. Metaphysical speculation is only of interest to philosophers and fantasy authors.

Quote:

Consciousness is scientifically undetectable. If scientists were not conscious as individuals then consciousness would not be thought to exist. There is no intra-subjective evidence that it does.

Absolutely false, unless you are using some strange definition of consciousness that is utterly different than what scientists are using, in which case the above is simply irrelevant.

Quote:

There is no scientific definition of consciousness that is generally accepted.

Generally accepted by whom? It comes as no surprise that the definition used by scientists is not accepted by many philosophers. Again, completely irrelevant. Scientists define consciousness in terms of various known phenomena

and processes, such as thought, perception, memory, and self-awareness. They leave out the metaphysical gibberish, because any definition which included metaphysics would be completely useless as a scientific definition.

Quote:

'What it is like' is the nearest thing to a working definition, and clearly, defined in this way, consciousness is a non-scientific substance. Crick suggests that we shouldn't even try to define it yet.

While armchair philosophers are wasting their time trying to figure out how to define a word, scientists are learning more and more about how the human mind works every day. Go figure.

Quote:

If I were James Randi I would put up a million dollars for the first neuroscientist to prove that consciousness arises from brains.

You would need to define consciousness first. Provide one which implies that consciousness is supernatural, and I am sure he would offer his million for it. Provide one which simply refers to those aspects of the mind which we know exist, and there would be no reason for him to offer the million, because it would not be a paranormal claim in any sense. On the contrary, it would be nothing more than what the current scientific evidence strongly indicates.

Quote:

Minds and brains are clearly causally connected in some way. But 'mind', conceived as a computational entity or process, has not yet been shown to be synonymous with consciousness, and many people assert that in the final analysis they are quite different things.

If you can't define what you mean by consciousness, then this entire line of argumentation is meaningless.

Quote:

Note that if this is the case then mind-brain/brain-mind causation becomes less of an intractable problem, since mind and matter are not all that there is and the interaction between the two may be mediated by something else.

That's a pretty big "if". Perhaps we should stick to trying to understand what we know is there, rather than speculating about what might be?

Quote:

If the scientific method was the only possibly way of researching consciousness then it would be sensible, as Metacristi suggests, to go along with the current best scientific theory and wait for it to improve. Unfortunately as yet there is no reasonable scientific theory of consciousness, and there won't be one until we can prove that consciousness exists.

Again, what do you mean by consciousness? I am getting the distinct impression that you are one of those people I mentioned who refuse to define what they mean by consciousness, instead only insisting that it is not any of those things scientists have managed to demonstrate to be brain processes.

Quote:

To disagree with Metacristi, there are existing non-scientific (or perhaps 'meta-scientific') explanations of consciousness that are perfectly plausible, so we do not necessarily have to wait for neuroscientists to figure it out before we do.

Are they testable? Or are they just unverifiable metaphysical speculation? If all you want is an explanation, and you are not concerned with knowing whether it is right, or actually being able to make any practical use of it, then I can provide you with a dozen or so. I'm just not going to lie to you and tell you that they are anything more than fantasies I (or somebody else) has dreamed up.

By Death Monkey (Kevin Dolan)

Monroe,

Quote:

I grasp the distinction just fine. You do not seem to grasp that it is not relevant. You are talking about

metaphysics, and I am talking about epistemology. Science describes things in terms of their physical effects. If you wish to claim that consciousness cannot be described in terms of its physical effects because it is supernatural, then you need to justify that claim. If you are claiming that such a description would be incomplete because it would leave out any non-physical components of consciousness, then you need to explain how you could possibly know that your, or anybody else's consciousness, has any non-physical components. You might as well claim that electrons have a non-physical properties. Maybe they do, but how would we ever know? And if we cannot know, why should we care?

Consciousness itself, as viewed by the owner, is not observable third-person.

This is simply not true. Any aspect of *your* consciousness which *you* are aware of having, has third-person observable effects on the world. The only question is whether or not it can be scientifically described in terms of those effects or not. In other words, whether or not they are supernatural effects.

Quote:

The physical effects on the brain and behavior are, but not the experience of consciousness itself.

Self-contradictory. Any aspect of your consciousness which you are aware of, affects your behavior, if not overtly, then at the very least at the brain-process level. And if you are not aware of it, then it is not your "experience of consciousness".

Quote:

But we do observe it from the first-person perspective.

And that affects your brain's behavior. End of story.

Quote:

The challenge for materialism is to show that the first-person perspective is explicable from a third-person physicalist standpoint.

The challenge of *science* with respect to *any* phenomena, is to provide a mechanistic explanation of it in terms of observable effects. Consciousness is no different than anything else in this respect. Appeals to metaphysical preconceptions are of no relevance to science.

Quote:

So what? This has absolutely no relevance to my points. None of this is even remotely incompatible with the hypothesis that your consciousness is something which your brain is doing.

It seems reasonable to say that consciousness is something your brain is "doing." (e.g., "causes") I just don't see how you can say it is one of the "physical" things it is doing, i.e. the stuff you can describe via neurobiology/physics/etc.

It is clearly physical. Anything which affects the world in some observable way, is physical. That is what "physical" means. You seem to be arguing against some sort of archaic metaphysical notion of "physical". Such notions have no place in modern science.

Quote:

In any event, this only means that it is impossible to prove that consciousness arises from matter, which is also irrelevant, because that is not what scientist try to do. Scientists do not need to prove that consciousness arises from matter. What they need to do is construct a scientific theory which describes how consciousness works.

If they are to reduce consciousness to material processes, how would this not be proof that it arises from matter?

That depends on what you mean by "arises". You seem to be talking about metaphysical reductionism. Science is concerned with epistemological reductionism. There is a very big difference. To see the difference, look at the example I already gave. It makes absolutely no difference to science whether the *metaphysical* fact is that minds created the physical world. As long as all detectable aspects of the mind can be *described* in terms of physical brain processes. Science does not make claims about ontological primacy, which seems to be what you are talking about.

Science attempts to describe the various features of the World in terms of our observations. Consciousness is clearly a feature of this World. Either it can be described scientifically, or it can not. Metaphysics simply does not enter into the question.

By Death Monkey (Kevin Dolan)

Monroe,

Quote:

There seems to be some confusion here. I am not disputing the claim that consciousness has effects on the body and brain. I am not even disputing the claim that you repeatedly bring up, that has no a priori certainty, that every difference in conscious experience makes some kind of difference with regard to the 3rd-person-observable aspects of the body and brain. Although this is not a priori certain, we have yet to find a case where different conscious experiences have no observable differences in the brain. Although I'm not familiar with much of the experimental data or the methodology, I'm willing to go with this claim. So basically we can establish a one-to-one correspondence between subjective experiences and brain events. But, as I noted before, the existence of a one-to-one correspondence between sets does not imply identity. I am arguing that consciousness, as it is experienced first-person and subjectively, not including the brain correlates, is not observable from a third-person standpoint.

Then you are either contradicting yourself, or do not understand what it means to observe something from a third-person perspective. What you are describing here is the classical philosophical problem of the distinction between the observable affects that something has, and the "thing in itself". The point is that this has no relevance to science. Absolutely none. Science does not claim that things are identical to the observable effects that our scientific descriptions use to describe them. Brain processes themselves are not *identical* to their physical effects. They are described in terms of them. Likewise, there is absolutely no reason to think that we should not be able to describe first-person experiences in terms of their effects. And if the scientific descriptions that result from such analysis indicates that in both cases it is the same phenomenon being described, then so be it.

Quote:

The only reason we ascribe mental states to other people is because that is the best explanation for their behavior; it allows us to interact with them very efficiently and make good predictions. But the skeptical problem of other minds remains. If you believe you have a solution to that problem, you should try to publish it; you will be famous!

No need. It is really very simple. I observe that my own consciousness affects my behavior. I observe that other people behave similarly. I observe that other people's behavior is caused by neuronal activity. I observe that I appear to be physically no different from other people. I therefore conclude that my own behavior is caused by my neuronal activity, and that my consciousness is somehow a function of that neuronal activity.

It isn't exactly rocket science. It is just a simple matter of looking at the evidence available, and drawing the obvious conclusions, no matter how counter-intuitive they may seem to be.

It is simple. If my consciousness affects my behavior, then there are exactly two options:

- 1) Other people also have consciousness.
- 2) Other people do not have it, and have other causes for their similar behavior.

The second option would indicate some physical difference between other people's brains, and my own. It would also imply that my own consciousness, whatever it is, is interacting with my special brain in some way.

Now which of these options is the more reasonable one?

Indeed, if I were to believe, as you seem to, that my own consciousness is not a brain process, but instead some supernatural agency interacting with my brain in some way, I would have to conclude that other people *did not have*

it, since their behavior appears to be completely determined by brain processes, and there is no indication of any supernatural influences. Everybody but me is a p-zombie.

Sounds kind of ego-centric to me. I'm not special. My behavior is determined by my brain processes, just like everybody else's. And since my consciousness clearly affects my behavior, the conclusion is obvious. My consciousness is a brain process, and other people have it too.

By Death Monkey (Kevin Dolan)

Monroe,

Quote:

Are you still not understanding what I am saying? You look out the window and see birds, trees, blue sky, etc. At the same time, I look at your brain and see neurons firing away. I don't observe your experiences by observing your brain.

I disagree. You are observing my experiences. What you are not doing is *experiencing* them.

Quote:

You basically seem to be saying that there are no subjective experiences, or that such things are "irrelevant to science."

I am not saying that at all. I am saying just the opposite. They do exist. They are relevant to science. They are physical. They are brain processes.

Quote:

I am not saying that there is some Kantian thing-in-itself which we don't have access to, I'm simply saying that subjective experiences are private, and brain processes are not.

They are only private in the sense that only you experience your experiences. That does not mean that I cannot observe them. I do not need to digest your dinner to observe the process of you digesting your dinner. Nor do I need to experience your experiences in order to observe the process of you experiencing them.

Quote:

As for whether it is reasonable to assume that subjective experience can be reduced to brain processes, descriptively and logically, this seems highly implausible to me without some kind of justification other than "it's reasonable to assume so."

Like I said before, the only alternative is that consciousness is some other "thing" which interacts with the brain in some way. Given the profound lack of evidence to support this possibility, the position that consciousness is reducible to brain processes is the most reasonable position to take.

Quote:

No need. It is really very simple. I observe that my own consciousness affects my behavior. I observe that other people behave similarly. I observe that other people's behavior is caused by neuronal activity. I observe that I appear to be physically no different from other people. I therefore conclude that my own behavior is caused by my neuronal activity, and that my consciousness is somehow a function of that neuronal activity.

This is the clearest thing (to me) that you've said so far. So you recognize that our access to our own consciousness is not simply by its "physical effects."

Of course. Our access to our own consciousness is a function of the fact that we *are* those processes which we call consciousness.

Quote:

To put it in one philosophical way, even if what we refer to with "consciousness" is the same as "brain", the words have different senses. Having different senses means that we use different properties of the thing to refer to it. For example, the morning star and evening star. The property of being the first visible heavenly body in the evening is

different from the property of being the last one in the morning. But it turns out that "that which appears first in the evening" and "that which appears last in the morning" pick out the same planet, Venus. In the case of consciousness, in one's own case one picks out whatever it is by the properties of subjective experience. If it is the same thing, it is picked out differently by referring to it as the brain process.

Yes, that is exactly my point. What I think of as my first-person consciousness, and the physical activity in my brain, are not *obviously* the same thing. I conclude that they are the same thing, based on the available evidence.

Quote:

Now, what we have seen while it is unclear whether there are two different referents picked out, there are definitely two different kinds of properties used in the senses. I, on the one hand, prefer to use "consciousness" to refer only to those properties of subjective experience, not to some substance which has those properties. What I am claiming is that you cannot deduce these properties from the physical properties of the brain.

The problem here is that your choice of terms presupposes your claim. You can define subjective properties to be only those properties which are properties of your subjective experiences, but to also stipulate that they are not properties of any substance (or physical process), is simply to presuppose that your subjective experiences are not just brain processes.

Dualism that only results from your choice of how to define your terms, is not really dualism. Simply defining the word "consciousness" to simultaneously refer to properties of subjective experiences, and to *not* refer to any properties of brain processes, does not magically make your subjective experiences not brain processes. It just means that your definition of consciousness is incompatible with the possibility of them being brain processes.

Quote:

I don't disagree that belief in other minds is the most reasonable option. However, the skeptical argument says only that we can't know for sure whether everyone else is a p-zombie.

We can't know anything about the external world for sure. That doesn't mean that there is no such thing as scientific knowledge about the external world. Just that such knowledge is never certain.

Quote:

Also, I disagree when you say, "The second option would indicate some physical difference between other people's brains, and my own." It is logically possible for everyone's brain to have the same physical properties without having the same subjective-experience properties. For example, p-zombies or qualia inversion.

Not if those subjective experiences have physical effects. As soon as you accept that consciousness is causally efficacious, you have rejected epiphenomenalism as a possibility. Qualia inversion and p-zombies then go out the window.

If p-zombie DM's brain activity is the only causal factor of his behavior, and I am physically identical to p-zombie DM, then it logically follows that my brain activity is the only causal factor of my behavior. If my consciousness is also a causal factor of my behavior, then it logically follows that either the p-zombie's behavior will be different from mine, or he will have some other brain activity I don't have, which provides the same causal influences on his behavior that my consciousness provides for me, which means that our brains must be different. Either way, the premise of being physically identical is violated.

By Death Monkey (Kevin Dolan)

Nomos,

Quote:

I disagree. You are observing my experiences. What you are not doing is experiencing them.

I think you mean that he is observing behaviour.

Since the scientific theory which I am advocating holds that experiences *are* a type of behavior (namely brain activity), there is no contradiction here. That is the whole point.

Quote:

It is impossible to observe experiences.

That is nothing more than the assertion that experiences are not any type of physical activity. Clearly I do not agree with this assertion.

Quote:

I know of no person except you who has ever suggested that it was. Among philosophers and scientists it is accepted that experiences are first-person things. There are no arguments about it.

On the contrary, many philosophers, and most scientists working in the field, agree that first-person experiences are brain processes. They simply avoid talking about "observing other people's experiences", because it is so easily confused with "experiencing other people's experiences". The meanings of those two phrases are very different, though.

Quote:

This is why consciousness cannot be shown to exist. Hence the 'other minds' problem.

There is no problem. I conclude the existence of other minds in exactly the same way I conclude the existence of anything else.

Quote:

It is why there is such a huge debate around the scientific validity of first-person reports of experience. First-person reports are not really scientific data, but without them there would be no scientific reason to believe that consciousness exists.

This is not true. The reason for the debate about the scientific validity of first-person reports is their unreliability. But this only applies to attempts to infer facts about what has happened to a person, from their experiences. It does not apply to the issue of determining the existence of their experiences.

Quote:

This is accepted by everyone in the field regardless of their preferred theory. In the end it seems no more than common sense.

This is simply false.

Quote:

I am not saying that at all. I am saying just the opposite. They do exist. They are relevant to science. They are physical. They are brain processes.

I think this is called 'identity dualism'. Mind and brain are different things with different names but are the same thing in fact.

I am not advocating identity dualism. I don't see why you feel the need to paraphrase my position. It is really extremely simple. My position is that subjective experiences are something which the brain *does*.

Quote:

It is usually considered incoherent. If your concept of an aeroplane is physical then how come it fits inside your brain?

Because it is a physical *process*, not a physical *object*. There is not "thing" inside of my brain which is my concept of an airplane. My concept of an airplane is a set of brain processes.

Quote:

How much does anger weigh?

How much does digestion weigh? Processes do not have weight, objects do. You are attacking the position that the mind is equivalent to the brain, or some object located inside of the brain. This is not my position.

Quote:

They are only private in the sense that only you experience your experiences. That does not mean that I cannot observe them. I do not need to digest your dinner to observe the process of you digesting your dinner. Nor do I need to experience your experiences in order to observe the process of you experiencing them.

So, what taste did you observe my dinner to have,

I did not experience eating your dinner. If I had observed in sufficient detail the brain processes that occurred while you ate, and knew enough about how your brain works, in principle I could tell you what sensation of taste you experienced.

Quote:

and did I like it?

Again, if I had made the necessary observations, and possessed sufficient knowledge of how your brain works, I could tell you whether you did or not. There is no need for me to experience your experiences in order to answer these questions.

Quote:

Consciousness is tricky topic and I think you may have made the mistake of meaning to put forward a scientifically acceptable view while in fact it contradicts science.

How does my view contradict science? I would be very surprised if it did, given that I work in a neuroscience lab, directly with scientists working in this field.

Quote:

If consciousness is causal then this would contradict the fundamental scientific assumption that the physical universe is causally complete.

Only if I were to make the ridiculous assumption that consciousness is non-physical. In any event, it is trivial to see that consciousness is causally efficacious. You may attach the label "consciousness" to something which you define to be causally inefficacious, but it is clear that when you do so, you are no longer talking about things like first-person experiences, because they clearly *do* affect our behavior. You are instead talking about something which, by definition, you could never possibly know about.

Quote:

According to science consciousness is not causal. It follows that observing someone's behaviour in order to infer the state of their consciousness is a complete waste of time, and a profoundly unscientific approach to the study of it. Perhaps the universe is not causally complete, but you'll have your work cut out arguing that with a physicalist.

This is complete nonsense. You are talking about epiphenomenalism, not physicalism. And epiphenomenalism is a decidedly non-scientific position. physicalism holds that consciousness *is* physical, and is, in fact, brain processes.

Quote:

Anyway, if observable behaviour is consciousness then what is it that we are trying to explain?

Consciousness. We want to explain how it works. Simply recognizing that it is a type of physical behavior (brain activity), is just the first step.

Quote:

Like I said before, the only alternative is that consciousness is some other "thing" which interacts with the brain in some way. Given the profound lack of evidence to support this possibility, the position that consciousness is reducible to brain processes is the most reasonable position to take.

I don't really want to disagree all the time, but this is not correct. The unfalsifiability of solipsism means that mind cannot be shown to reduce to brain. This does not prove that it does not, but if it does nobody will ever know it.

By this reasoning, nothing is knowable. Solipsism, like all unfalsifiable hypotheses, simply has no relevance to scientific epistemology.

Quote:

There is so much evidence that mind does not reduce to brain that it is overwhelming.

Name one piece of evidence. Note that I said evidence, not intuitive arguments or metaphysical presumptions.

Quote:

All that prevents the acceptance of it as a fact is a fear of its consequences for materialism. Materialists do not like their paradigm to be rattled any more than do theists.

This is complete garbage. Many materialists (myself included) arrived at their position *in spite* of a strong emotional desire to believe otherwise. I don't know where you got the idea that people *want* to believe that there is nothing more to the mind than brain activity. I don't know about you, but the thought of my own mortality scares the crap out of me. The thought that people I loved who are now dead, simply do not exist any more, depresses the crap out of me. I would *love* to believe that there is something more to the mind than just brain processes. I would love to be convinced that there is some sort of afterlife, or cosmic consciousness, or reincarnation, or pretty much any thing other than complete non-existence. But the evidence simply indicates otherwise. What I want has no bearing on how things really are.

Quote:

However there are signs that the irreducibility of mind to brain will soon be accepted, for something has to give way in consciousness studies before long. The literature is filling up with devastating criticisms of physicalism.

And yet, oddly, none of this seems to be slowing down the progress of scientific research into the mind. Go figure. Why don't you tell me where you think the current methodology of scientifically studying the mind will fail?

Quote:

Importantly, it is becoming increasingly apparent that mind and consciousness are not the same thing. This gives us three terms to work with, consciousness, mind and brain, and thus allows for a different kind of approach to the scientific 'problem of consciousness'. This is worth considering.

Apparent to whom, and in what way? How do mind and consciousness differ? It seems to me that you are just taking the old dualistic tactic of taking any aspect of the mind which we manage to explain in terms of brain processes, and saying "that's not consciousness. Consciousness is that other stuff that hasn't been explained yet".

Quote:

Of course. Our access to our own consciousness is a function of the fact that we are those processes which we call consciousness.

Here I agree. Other people's experiences are inaccessible to other people, and therefore not observable or measurable.

This does not logically follow. All it really means is that people's experiences cannot be *experienced* by other people. That does not mean that they are inaccessible, or that they are non-observable or unmeasurable. Those are epistemological assumptions *you* have made, and ones which are flatly contradicted by scientific evidence.

Quote:

Yes, that is exactly my point. What I think of as my first-person consciousness, and the physical activity in my brain, are not obviously the same thing. I conclude that they are the same thing, based on the available evidence.

I should trust your common sense on this one, which is what makes it non-obvious to you that mind and brain are the same thing.

Not common sense, intuition, which is notoriously unreliable.

Quote:

There is no evidence that they are.

This is simply false. You are completely ignoring the existence of an entire field of scientific research, and all of the progress that has been made in that field over the past several decades.

Quote:

The problem here is that your choice of terms presupposes your claim. You can define subjective properties to be only those properties which are properties of your subjective experiences, but to also stipulate that they are not

properties of any substance (or physical process), is simply to presuppose that your subjective experiences are not just brain processes.

Well, I would have thought that it was fair enough for Monroe to claim that subjective experiences are no more or less than subjective experiences. It is not necessary to presuppose that they are not caused by physical process in order to conclude that subjective experiences are not brain processes.

He did not conclude that they are not brain processes. He assumed that they were not, by defining a term in such a way that would render it self-contradictory if they are. If you think this is ok, then that's fine. You are just simply wrong.

Quote:

Further evidence is the fact that we do not call subjective experiences brain processes, we use different terms.

This is evidence of nothing more than the beliefs of the people who created the language we use.

Quote:

When we look into a brain we find no sign of subjective experiences. As neurophysiologist Karl Pribram says, looking for consciousness in the brain is about as useful as digging to the centre of the earth in search of gravity.

I won't find computation by looking inside a microchip, either.

Quote:

Unfortunately you, no more than anyone else, can solve the problem of consciousness by redefining 'subjective experience' as 'brain process'. Evidence is required, and a working theory of some sort.

Both of which we have. Like I said, you are either completely ignoring an entire field of scientific research, or you have taken the epiphenomenalist approach of defining consciousness to refer to something which we could never even know we have (ie something causally inefficacious).

Quote:

We can't know anything about the external world for sure. That doesn't mean that there is no such thing as scientific knowledge about the external world. Just that such knowledge is never certain.

That's true. But don't forget that consciousness is the one thing that we can know anything about with certainty.

Your own consciousness is the only thing you can know anything about with any certainty, and the only thing which you can know about it with any certainty, is that it exists. This in no way support your position that consciousness is not scientifically knowable.

Quote:

Not if those subjective experiences have physical effects. As soon as you accept that consciousness is causally efficacious, you have rejected epiphenomenalism as a possibility.

Epiphenomenalism actually makes more sense if consciousness is causal, since it would mean that mind-brain interaction is just interaction between emergent levels of complexity in a system, which we already know occurs both upwards and downwards in emergent and hierarchic dynamic systems.

But then that is no longer epiphenomenalism. That is just exactly what physicalism claims! The various properties which we call "consciousness" are emergent properties of the complex dynamical behavior of the brain.

Epiphenomenalism differs from physicalism in the claim that these emergent properties are not logically implied by, nor reducible to, the brain activity which they are emergent properties of. I am not sure what you think physicalism is, but I suspect that bears little resemblance to what it actually claims.

Quote:

However epiphenomenalism makes little sense if consciousness is non-causal, since it would assert that brain can cause mind but mind cannot cause brain. This would be the only example of one-way causation in the known universe. It is known as 'assymetric supervenience' in the trade, and this is not generally considered very plausible.

I agree. So can we dismiss this nonsense about non-causally efficacious consciousness?

Quote:

Mind-brain/brain-mind interaction becomes less of an intractable problem to wrestle with if you conjecture that there is a third term involved.

Which gives exactly the point I raised before. If it is not the case that consciousness is just brain activity, then it follows that whatever consciousness is, it interacts with the brain in some way. These interactions should be detectable. This means that either we can study consciousness in the same way that we study anything else (by observing how it interacts with other things), or that consciousness is supernatural.

In either case, I will give serious consideration to such a hypothesis when there is evidence to support it. Discovering some of these mysterious interactions would be a good start. So far, there is absolutely no indication of any such third-party interacting with the brain.

By Death Monkey (Kevin Dolan)

Monroe,

Quote:

When you observe my digestion of my dinner, the third-person observation will tell you all the facts about my digestive system. But when you observe my brain while I am looking at a picture, you will not know from this what my experience is like, what kinds of beliefs I am having, and such.

I will not know what your experience is like, because the phrase "know what your experience is like" is just another way of saying "experience your experience". This is not an example of a fact about your experience. As for the others, I see no reason why third-person observations of your brain processes, combined with a thorough knowledge of how your brain works, would not tell me what kinds of beliefs you are having.

Quote:

The extent to which you can conclude what thoughts and qualia I am having by observing my brain is only by establishing correlations experimentally. However, a logical deduction is not possible, which is required by the thesis that consciousness is reducible to the physical aspects of the brain.

All scientific conclusions are drawn by establishing correlations experimentally. All that the thesis that consciousness is reducible to brain activity requires, is that I should be able (in principle) to use a *complete* scientific theory about how consciousness works, to deduce all facts about your consciousness from my observations of your brain processes.

Quote:

Like I said before, the only alternative is that consciousness is some other "thing" which interacts with the brain in some way. Given the profound lack of evidence to support this possibility, the position that consciousness is reducible to brain processes is the most reasonable position to take.

Or property dualism, which holds that there are properties of the brain which are invisible to outside observation, and visible to the brain itself. These properties are not reducible to the 3rd-person observable properties. In case this is unclear, it's like the fact that electric charge is not reducible to mass.

This is self-contradictory. If they are "visible" to the brain, then they must affect the brain in some way. if they do, then those effects are 3rd-person observable. In this case, you are back to exactly the scenario I indicated above.

If these properties are not 3rd-person observable, they cannot have any affect on your brain activity. If they have no affect on your brain activity, then they are not "visible to the brain itself".

Quote:

I define subjective properties to be that which we consciously experience. I am not saying that they are not

properties of some substance. I am just saying that my term does not refer to the substance, but to those properties only. My claim that they are not reducible to physical brain processes is not part of the definition, but a substantive claim. Now, obviously the definition of 'subjective experience' has a different sense than that of 'brain process.' So clearly they pick out, at the very least, different properties of something. However, as we see in mathematics, sometimes things be defined with different senses, but it turns out that they are equivalent. For materialism to work, it must show that they are logically equivalent notions, even if defined in different senses.

Actually, all materialism (or more accurately, science) needs to do, is construct a falsifiable mechanistic theory for how it works. Logical proofs don't enter into it. Scientific theories cannot be logically proven to be true, which seems to be what you are demanding. They can only be shown to consistent with all available data, and to be able to accurately predict new observations.

Quote:

Not if those subjective experiences have physical effects. As soon as you accept that consciousness is causally efficacious, you have rejected epiphenomenalism as a possibility. Qualia inversion and p-zombies then go out the window.

The nature of qualia inversion is such that there wouldn't be any different effects on behavior. They still cause things, just the same things. Whether it is logically possible for qualia inversion to happen without any difference in the brain seems obvious, unless you can show that brain processes logically imply certain qualia. Same with p-zombies.

With respect to qualia inversion, if two different qualia both had *exactly* the same physical effects on your brain, you would not know that they were different. So in what sense is it meaningful to say that they *are* different? The very definition of "qualia" defines them in terms of your subjective experience. If you can't tell the difference between two experiences, then in what sense is it meaningful to say that they are different?

As for p-zombies, your above statement is simply false. A p-zombie's behavior is *completely* determined by his brain processes. If I have the same brain processes as the p-zombie, then it logically follows that my behavior is also completely determined by my brain activity. You can assert the existence of something else, which the p-zombie does not have, but it is meaningless to say that it has caused my behavior, when my behavior would have been exactly the same had it not been there.

Quote:

If p-zombie DM's brain activity is the only causal factor of his behavior, and I am physically identical to p-zombie DM, then it logically follows that my brain activity is the only causal factor of my behavior. If my consciousness is also a causal factor of my behavior, then it logically follows that either the p-zombie's behavior will be different from mine, or he will have some other brain activity I don't have, which provides the same causal influences on his behavior that my consciousness provides for me, which means that our brains must be different. Either way, the premise of being physically identical is violated.

None of this "logically follows" unless you assume the mind-brain identity thesis.

Wrong. The only premises required are those I stated:

Premise 1) The p-zombie's brain activity is the only causal factor of its behavior.

Premise 2) I am physically identical to the p-zombie.

Premise 3) My consciousness is a causal factor in my behavior.

The first two premises are required by the definition of what a p-zombie is. The third premise is clearly not compatible with the first two. At least one of them must be wrong. Either p-zombies are logically impossible, or my consciousness is not a causal factor in my behavior. I know that the latter is false, so I conclude that p-zombies are not possible.

Quote:

P-zombie DM could have different causes for the same effects, or his brain could just operate in a way that matches yours, except that the causal level is deleted for a lot of things. (like the earth continuing to have the same orbit without the sun being there... its not physically, but logically possible, since we could describe such motion mathematically without contradiction.)

What you are essentially saying is that it is logically possible that our scientific conclusions about how the brain works could all be wrong. This is certainly true. It could be supernatural, as you suggest. My point is that epiphenomenalism is not logically consistent with our current scientific knowledge. If you wish to reject science in order to maintain a belief in property dualism, go right ahead.
By Death Monkey (Kevin Dolan)

Monroe,

Quote:

I will not know what your experience is like, because the phrase "know what your experience is like" is just another way of saying "experience your experience". This is not an example of a fact about your experience.

Some reasons why knowing what it's like is a fact, and is not the same as experiencing something:

1) I know constantly what it is like to see red. I don't experience red constantly.

Is that really true? Do really always know what it is like to see red, even when you aren't thinking about it? I think that it is very important to be very precise in language here. I would say that your *knowledge* of what it is like to see red is your memory of having experienced seeing red, and that *knowing* what it is like to see red is you being aware of your experience of seeing red, or remembering having seen red before.

In the former case, it is not a fact which you "know". It is a memory of an experience, which is a physical configuration in your brain. In the latter case, it is not a fact which you know, it is an experience you are having.

Quote:

2) If I have reasons to believe that your experience is the same as my experience, then I know what your experience is like, without experiencing it.

In this case, the only "fact" involved is the fact that my experiences are the same as yours. What it is like for either of us to see red, is not a fact. That my experience of seeing red is, or is not, the same as yours, is a fact.

Quote:

3) Any particular experience is a real phenomenon in the universe. To not know anything about this phenomenon is to leave out some facts of the universe. If you do not know what is going on in one of these phenomena, you are failing to describe a real thing, not merely failing to be experience

This has no relevance to my argument. My point is that I can know all of the facts about your experience, without having that experience myself. You have argued that I will not know what it is like for you to see red. My point is that what it is like for you to see red is not a fact about your experience. It just sounds like it should be, because of the way we refer to it.

Quote:

As for the others, I see no reason why third-person observations of your brain processes, combined with a thorough knowledge of how your brain works, would not tell me what kinds of beliefs you are having.

And I see no reason why they would! Can you give an example of how you can deduce a certain belief-state from a certain brain process? The absence of knowledge of such examples would mean there is no evidence for the position that they exist. Of course, absence of evidence does not prove absence, but we are talking about scientific, probabilistic epistemology here.

There are many, many examples of this in modern neuroscience. Of course, we are no where near the point of being able to electronically read minds, but we can certainly draw reliable conclusions about subjective states based on looking at brain activity.

Quote:

All scientific conclusions are drawn by establishing correlations experimentally. All that the thesis that consciousness is reducible to brain activity requires, is that I should be able (in principle) to use a complete scientific theory about how consciousness works, to deduce all facts about your consciousness from my observations of your brain processes.

If you're admitting that there is no way to deduce consciousness from the brain without establishing experimental correlations between brain states and reports of subjects, then you are admitting I am right. Because that is simply all I am arguing.

Then you are arguing against a strawman version of the scientific mind-brain hypothesis. Like I said, all scientific conclusions are drawn this way.

Quote:

In case that was unclear, sorry, but that's what I meant this whole time, just that no a priori connection between the brain and consciousness exists.

Again, you are talking metaphysics here, not science.

Quote:

This is self-contradictory. If they are "visible" to the brain, then they must affect the brain in some way. If they do, then those effects are 3rd-person observable. In this case, you are back to exactly the scenario I indicated above.

If these properties are not 3rd-person observable, they cannot have any affect on your brain activity. If they have no affect on your brain activity, then they are not "visible to the brain itself".

It was just a manner of speaking, but unfortunately it was unclear. I really meant that the properties are "visible" to themselves, or visible to the mind or consciousness (which is what they collectively are). Anyway, yes I do think that these properties have effects on the 'physical' properties, but this doesn't mean that are their effects.

This brings us back to what I said before. Either consciousness is a brain process, or it is something else which interacts with the brain. If the latter, either it can be studied scientifically by looking at these interactions, or it is supernatural. There is no evidence for any supernatural interactions in the brain.

Quote:

Actually, all materialism (or more accurately, science) needs to do, is construct a falsifiable mechanistic theory for how it works. Logical proofs don't enter into it. Scientific theories cannot be logically proven to be true, which seems to be what you are demanding. They can only be shown to consistent with all available data, and to be able to accurately predict new observations.

I'm not sure what you mean by a "mechanistic theory for how it works." I think I have shown that the issue of mind-brain identity rests on whether a logical reduction of mind to brain can be found.

By a mechanistic theory for how it works, I mean a theory which explains all of the various phenomena which we lump together with the term "consciousness", in terms of observable physical processes.

Quote:

With respect to qualia inversion, if two different qualia both had exactly the same physical effects on your brain, you would not know that they were different. So in what sense is it meaningful to say that they are different? The very definition of "qualia" defines them in terms of your subjective experience. If you can't tell the difference between two experiences, then in what sense is it meaningful to say that they are different?

If my color qualia are inverted with respect to yours, you do not know it, and I do not know it, because we do not have access to each others' consciousnesses. But it is meaningful, in that when I experience blue, you experience yellow, but we call them the same names and behave as if they were the same experience. If you're saying that this is not meaningful, then you are bringing up much deeper philosophy of language issues which I cannot enter into at this time. Suffice it to say that positivism has quite a bad rep these days.

I would say that it would only be relevant if our brains were physically identical, but they are not. I think that it is ridiculous to imagine that *any* of my qualia are going to be *exactly* like any of yours.

Quote:

As for p-zombies, your above statement is simply false. A p-zombie's behavior is completely determined by his brain processes. If I have the same brain processes as the p-zombie, then it logically follows that my behavior is also completely determined by my brain activity. You can assert the existence of something else, which the p-zombie does not have, but it is meaningless to say that it has caused my behavior, when my behavior would have been exactly the same had it not been there.

It does not logically follow. The stipulation is that the p-zombie's brain and behavior are identical to yours. What the causal structure of the p-zombie is is irrelevant. It only needs to track the normal human in physical properties exactly. Remember my example of a planet that goes in the same orbit without the sun being there to cause it; the p-zombie is like this, and it is logically consistent as I have already argued. At any rate, we could simply stipulate that the causal structure of the p-zombie is different from the normal human.

Again, if you are going to reject the notion of naturalism, the entire thing goes out the window. If it all works magically anyway, then anything is possible. Maybe little invisible gremlins that live in the center of the moon are actually controlling everything, and are just doing so in such a way as to give us the impression of physical causality, because it amuses them to do so?

The point is that if we can explain the behavior in terms of brain processes, it is silly to assume that the brain processes are not actually the cause, and that instead some undetectable spirit is.

Quote:

What you are essentially saying is that it is logically possible that our scientific conclusions about how the brain works could all be wrong. This is certainly true. It could be supernatural, as you suggest. My point is that epiphenomenalism is not logically consistent with our current scientific knowledge. If you wish to reject science in order to maintain a belief in property dualism, go right ahead.

This does not make any sense. I am talking about logical possibilities. I am not rejecting what science shows to be actually (and contingently) the case. I am not saying that for all we know, science could be all wrong. No, I am saying that it is logically consistent for things to have been different than they are. The issue of whether p-zombies are logically possible is different from whether any p-zombies actually exist.

And also utterly irrelevant. What is relevant is whether or not *our* behavior is caused by brain activity. Science indicates that it is. This means one of three things:

- 1) Science is wrong, and our behavior is really caused by something else.
- 2) Consciousness is brain activity.
- 3) Consciousness is not brain activity, and is not causally efficacious.

I reject option (1) because it is vacuous supernaturalism. I reject option (3) because what *I* think of as "consciousness" clearly *is* causally efficacious, and therefore whatever it is that is not causally efficacious, which is being called "consciousness", is not something which I have any reason to believe I actually have.

Remember, the issue here is not whether it is *possible* for there to be some hypothetical world where brain activity does not cause our behavior, but instead spirits do. The issue here is whether or not *our* world is such a world.

Science indicates that our behavior is caused by brain activity. We can either accept science, and reject the superstitious notion of spirits, or we can reject science, and believe any logically self-consistent hypothesis which makes us feel good.

By Death Monkey (Kevin Dolan)

Nomos,

Quote:

Since the scientific theory which I am advocating holds that experiences are a type of behavior (namely brain activity), there is no contradiction here. That is the whole point.

No contradiction? It is one thing to say the experiences are caused by processes. It is quite another to say that they are brain processes. The latter makes no sense. If experiences were brain processes then we would not be able to have them. When was the last time you observed your brain?

I don't follow your reasoning here at all. I think one problem is that you are using the term "observe" in a very generic way. Exactly what do you mean when you ask me about observing my own brain? Since my thoughts are brain processes, and my memories are stored as patterns in my brain, I would say that every time I think about anything, I am observing my brain.

Quote:

That is nothing more than the assertion that experiences are not any type of physical activity. Clearly I do not agree with this assertion.

No, I am not saying that experiences are not connected with brain processes. But they are not a type of brain process. I can observe your brain processes, but not your experiences. This makes it clear that they are not the same thing, anymore than a steam-whistle is the same thing as steam.

Where is your evidence that experiences cannot be observed. Note again that observing your experiences is **not** the same thing as experiencing them.

Quote:

On the contrary, many philosophers, and most scientists working in the field, agree that first-person experiences are brain processes.

This is not the case. A small (and dwindling) number of them argue this. Most assume that brain processes cause experiences, but that experiences are not identical to brain processes. They also all agree that both views are just conjectures.

Apparently none of the neuroscientists or psychologists I work with got the memo.

Quote:

They simply avoid talking about "observing other people's experiences", because it is so easily confused with "experiencing other people's experiences".

They avoid saying that because it would be a daft thing to say. Nobody has ever claimed that experiences can be observed. Behaviour can be observed, that is all. This is why Behaviourism exists as a doctrine, or a methodology.

Could you do me a favor and tell me right now if all of your arguments are going to simply amount to assuming that my position is false, and the concluding from that assumption that they are, so that I will know whether I am completely wasting my time here? I just said that my position is that experiences are a type of behavior. You just mentioned behaviorism, well, that is exactly what behaviorism says.

Quote:

There is no problem. I conclude the existence of other minds in exactly the same way I conclude the existence of anything else.

Sorry, but no you don't. You infer the existence of consciousness in others from their behaviour, but you do infer the existence of carpets from their behaviour.

Actually, I do. If carpets did not interact with anything else, I would have no idea that they exist. That is behavior.

Quote:

This is not true. The reason for the debate about the scientific validity of first-person reports is their unreliability. But this only applies to attempts to infer facts about what has happened to a person, from their experiences. It does not apply to the issue of determining the existence of their experiences.

I'm afraid it does. I agree that it's a bit pedantic to argue this, since common sense tells us that other people are conscious, but scientists and philosophers are notoriously pedantic.

I know. I happen to be a scientist, working in the field of neuroscience. I am well aware of the difficulties involved in scientific research concerning subjectivity.

Quote:

I am not advocating identity dualism. I don't see why you feel the need to paraphrase my position. It is really extremely simple. My position is that subjective experiences are something which the brain does.

Your position, which you've stated quite clearly, is that brain and mind are the same thing. Is this not what you think?

This is becoming quite tedious. I just told you what my position is. My position is that subjective experiences are something which the brain does. If you want to attack my position, then deal with what I said. The statement "mind and brain are the same thing" is far too imprecise and vague for me to either agree, or disagree with.

Quote:

Because it is a physical process, not a physical object. There is not "thing" inside of my brain which is my concept of an airplane. My concept of an airplane is a set of brain processes.

A brain process is a brain process. Such a process may cause an experience, but cannot be an experience.

That is what you believe. I think you are wrong.

Quote:

If it were then we would not be able to observe brain processes.

This does not logically follow at all.

Quote:

I did not experience eating your dinner. If I had observed in sufficient detail the brain processes that occurred while you ate, and knew enough about how your brain works, in principle I could tell you what sensation of taste you experienced.

Again, if I had made the necessary observations, and possessed sufficient knowledge of how your brain works, I could tell you whether you did or not. There is no need for me to experience your experiences in order to answer these questions

I think you need to think about that a bit further. It is very obviously not the case. If you search around you will not find one researcher who agrees with you on this, and there's a reason why they don't.

I am afraid your attempted appeal to authority falls rather flat on me. I happen to work with several researchers who do agree with me on this.

Quote:

How does my view contradict science? I would be very surprised if it did, given that I work in a neuroscience lab, directly with scientists working in this field.

You assume that consciousness causes behaviour. This contradicts the idea that the physical universe is causally complete.

This argument assumes that consciousness is not a physical process. I have already explained that I think that it is.

Quote:

There is no scientific evidence that freewill exists, and most scientists argue that it cannot exist.

Who said anything about free-will?

Quote:

Only if I were to make the ridiculous assumption that consciousness is non-physical.

Can you explain why you find this assumption ridiculous? If you have a good argument you ought to write a paper. It would revolutionise the field if you could make a good case.

I'm afraid such a paper would hardly be considered new. This is the basic position upon which *all* scientific research into consciousness is based. After all, science can only study physical phenomena.

Quote:

In any event, it is trivial to see that consciousness is causally efficacious.

Same response. It would revolutionise science and philosophy if you could make a good case.

No, it wouldn't. Clearly the fact that we are *physically* discussing our own consciousnesses right now, demonstrates that our own consciousnesses have had some effects on our brains. Clearly whatever it is that you think of as consciousness has some effect on your brain. If it did not, your brain would not be able to think about it.

Quote:

And epiphenomenalism is a decidedly non-scientific position. physicalism holds that consciousness is physical, and is, in fact, brain processes.

How can an experience be physical? It may be caused by a brain process, the outcome of a process, but it cannot be a process. The idea is incoherent.

How so? You keep saying this, but have offered no explanation why.

Quote:

By this reasoning, nothing is knowable. Solipsism, like all unfalsifiable hypotheses, simply has no relevance to scientific epistemology.

If you read the academic literature on consciousness you will find a great deal of discussion about solipsism, idealism, physicalism, and so forth. They are directly relevant to the topic.

Sorry, I have read lots of scientific literature on the subject, but you seem to be talking about philosophical literature on the subject. Specifically, metaphysical literature. Metaphysics has absolutely no relevance to science.

Quote:

If we could show that consciousness is a brain process then we would have shown that solipsism is false.

That depends on what you mean by "show". If I can show that anything other than my own mind exists, then I have shown solipsism to be false. No amount of scientific evidence can ever *prove* this to be true. That does not mean that there cannot be scientific evidence to support the hypothesis that it is.

Quote:

As we cannot show that it is false we cannot show that consciousness is a brain process. Not ever, not if we have ten centuries to do it in. You can draw varying conclusions from this, but you cannot argue that it is not a relevant fact.

It is a completely irrelevant fact to scientific epistemology, because what you are talking about is absolute proof, which has no place in scientific epistemology. I cannot provide absolute proof for *any* claim about reality.

Quote:

Name one piece of evidence. Note that I said evidence, not intuitive arguments or metaphysical presumptions.

The strongest evidence is the inability of scientists to come up with a single working theory of how mind arises from matter.

What do you mean by working theory? If you mean scientific theories which partially explain what is going on, you are dead wrong, because we have them. If you mean a complete theory which explains how consciousness arises from neuronal activity, then the absence of such a theory is **not** strong evidence that it doesn't. Before we could even *attempt* to construct such a theory, we would have to understand, in detail, how the brain works, and we simply do not yet.

Quote:

Another is the persistent conclusion of people who explore consciousness, known as the perennial philosophy. These people assert that they know that consciousness is, ultimately, not mind or matter. They may all be mis-reporting their first-person experiences of course, but you'd have to be a very paranoid conspiracy theorist to think that they were all lying, and if it's a coincidence then it's a quite astonishing and inexplicable one.

Gee, it never occurred to you that maybe they have just reached the most intuitively compelling conclusion? I'm sorry, but I don't place much faith in intuition, particularly when it contradicts scientific evidence.

Quote:

This is complete garbage. Many materialists (myself included) arrived at their position in spite of a strong emotional desire to believe otherwise. I don't know where you got the idea that people want to believe that there is nothing more to the mind than brain activity. I don't know about you, but the thought of my own mortality scares the crap out of me. The thought that people I loved who are now dead, simply do not exist any more, depresses the crap out of me. I would love to believe that there is something more to the mind than just brain processes. I would love to be convinced that there is some sort of afterlife, or cosmic consciousness, or reincarnation, or pretty much any thing other than complete non-existence. But the evidence simply indicates otherwise. What I want has no bearing on how things really are.

Well, it's not very helpful to start calling people's arguments garbage, not unless you can counter them.

I did not call your argument garbage. I said that your attempt to dismiss my position by saying that the people who hold it only do so because they want it to be true, is garbage. More specifically, it is an ad-hominem attack. And I did counter it. I countered it with a specific counter-example.

Quote:

But other than that this is well said. I realise that you believe you are being intellectually rigorous by being a materialist, and I respect that. However you should bear in mind that materialism is a metaphysical doctrine, not a scientific theory.

Stop right there. I am not a metaphysical materialist. I do not hold any metaphysical beliefs. I consider metaphysics itself to be completely vacuous. My position is an epistemological one. I believe that consciousness can be scientifically understood in exactly the same way as any other phenomenon. That is an epistemological position, not a metaphysical one. It is also a scientific theory.

Quote:

There is no evidence that it is true, none at all. I'm sure you will not believe this, but try finding any conclusive scientific evidence for materialism. There isn't any. It is well known that materialism is no more than a metaphysical conjecture.

Which is why I reject it, right along with dualism, idealism, solipsism, and all other metaphysical speculation. I am not sure where you got the idea that I was arguing for a metaphysical position, but it certainly was not from anything which I actually *said*.

Quote:

And yet, oddly, none of this seems to be slowing down the progress of scientific research into the mind. Go figure. Why don't you tell me where you think the current methodology of scientifically studying the mind will fail?

Scientific researchers have yet to make any progress on the origins of consciousness. So far they have not even been able to define it. The current methodology fails because consciousness is not scientifically observable, and therefore has to be researched in a first-person way.

This is simply false. You clearly are either ignorant of the scientific research in this field, or have conveniently redefined "consciousness" to not include any of the stuff which scientists have been gaining understanding of over the past few decades.

Quote:

How do mind and consciousness differ?

In my opinion mind and matter reduce to consciousness.

It seems to me that you are just taking the old dualistic tactic of taking any aspect of the mind which we manage to explain in terms of brain processes, and saying "that's not consciousness".

Yes, that's exactly what I'm doing. But it's not dualistic. Materialism is dualism, so is idealism, both of which fail because of it.

What??? Materialism and idealism are both monisms.

Quote:

Dualism, cosmologically speaking, makes no sense. But it is not dualism to say that mind and matter are complementary substances, and reduce to something else that is not intrinsically dual.

That is just another form of monism. So what? I am not talking about metaphysics here. Science says nothing about metaphysics. Scientific theories will not tell us whether monism, dualism, materialism, idealism, or any other metaphysical position, is correct. What science can do, or at least try to do, is explain consciousness using scientific epistemology. So far, it is doing a very good job of this, and so far, the indication is that consciousness is something which our brains do.

Quote:

This does not logically follow. All it really means is that people's experiences cannot be experienced by other people. That does not mean that they are inaccessible, or that they are non-observable or unmeasurable. Those are epistemological assumptions you have made, and ones which are flatly contradicted by scientific evidence.

This is self-contradictory. Experiences can be experienced, but never observed from a separate viewpoint to that of the experienter. This is not an epistemological assumption but, as you say in your first sentence, just a bald fact.

OK. Prove it. Needless to say, I disagree. I don't think it is a fact at all.

Quote:

Not common sense, intuition, which is notoriously unreliable.

Every way we can reason is unreliable, or can be done badly, but common sense is less prone to error than complex scientific theorising, as is illustrated by the idiocy, now recognised, of Behaviourism, which contradicts

common sense and experience.

Whatever. I am a scientist. I know that intuition, while often useful, are very unreliable. I also know that scientific methodology is demonstrably reliable. You can believe whatever you want, but I'm not going to dismiss the scientific evidence just because I find it counter-intuitive.

Quote:

This is simply false. You are completely ignoring the existence of an entire field of scientific research, and all of the progress that has been made in that field over the past several decades.

What progress?

Go find a book on modern neuroscience. Any book should do.

Quote:

He did not conclude that they are not brain processes. He assumed that they were not, by defining a term in such a way that would render it self-contradictory if they are. If you think this is ok, then that's fine. You are just simply wrong.

Why is it ok to assume that consciousness is a brain process, but not ok to assume that it isn't?

It's not. I don't assume that consciousness is a brain process. That would be silly. I have concluded that it is, based on the scientific evidence.

Quote:

I won't find computation by looking inside a microchip, either.

Not exactly no, computation is a process, not a thing.

So is consciousness.

Quote:

But you will find that you can explain the output of the chip in terms of its inputs and its physical structure. This cannot be done for brains.

How do you know? So far it seems to be working quite well. Why do you assume that the brain is supernatural?

Quote:

Your own consciousness is the only thing you can know anything about with any certainty, and the only thing which you can know about it with any certainty, is that it exists. This in no way support your position that consciousness is not scientifically knowable.

Most philosophers now argue that consciousness is not scientifically knowable, along with many scientists. It's pretty obvious really. It's the reason science cannot explain it.

I am not interested in arguments from popularity. I could not care less what most philosophers believe about it. As for the scientists, in spite of your claims of impossibility, they seem to be doing a pretty good job of it.

Quote:

Epiphenomenalism differs from physicalism in the claim that these emergent properties are not logically implied by, nor reducible to, the brain activity which they are emergent properties of. I am not sure what you think physicalism is, but I suspect that bears little resemblance to what it actually claims.

Epiphenomenalism, in its usual guise, asserts that mind is an epiphenomenon, a waste product of the brain and non-causal. Physicalism states that conscious experiences are made out of matter.

No, physicalism states that conscious experiences are physical processes. i have never heard any physicalist claim that conscious experiences are physical objects.

Quote:

So can we dismiss this nonsense about non-causally efficacious consciousness?

Well, I think that it's nonsense, but scientists think it's true, so we can't just dismiss it out of hand, unfortunately.

I have never met a scientist who works in the field of consciousness, who thinks that consciousness is causally inefficacious. That would be insane, since it would imply that consciousness cannot be studied scientifically. I have also never met a neuroscientist who thinks that consciousness is causally-inefficacious, although I guess there could be some out there. I am curious as to where you are getting all of these ideas about what scientists in the field do and do not believe, from?

Quote:

If it is not the case that consciousness is just brain activity, then it follows that whatever consciousness is, it interacts with the brain in some way. These interactions should be detectable.

I agree. It's very odd that we cannot find these interactions. Perhaps there's a fault in our assumptions.

In your assumptions, not mine. I do not assume that consciousness is not just brain activity, so the lack of these interactions is exactly what I expect. it is your belief that consciousness is not brain activity, and that it is causally efficacious, which implies that these interactions should be there.

Quote:

This means that either we can study consciousness in the same way that we study anything else (by observing how it interacts with other things), or that consciousness is supernatural.

Not supernatural, just meta-scientific. These are quite different things. There cannot be anything supernatural about consciousness, it's part of the natural world. Why it evolved and what survival purpose it serves is a scientific mystery, but it must must be a naturally occurring 'substance'.

There is no difference. Either it functions according to natural laws, which can be inferred from observation, or it doesn't. If it doesn't, then why it doesn't is unknowable to us, and therefore irrelevant. You have no way of knowing whether consciousness is just "meta-scientific", as you put it, or does just not conform to any natural laws at all. Your assertion that "it must be a naturally occurring substance", is pure conjecture.

Quote:

In either case, I will give serious consideration to such a hypothesis when there is evidence to support it. Discovering some of these mysterious interactions would be a good start. So far, there is absolutely no indication of any such third-party interacting with the brain.

The evidence is there. You just don't accept it as evidence. Have you considered that there might be a reason why consciousness, and brain-mind/mind-brain interaction, seems to be an intractable scientific mystery?

It doesn't seem very intractable to me. We are making progress on it all the time.

By Death Monkey (Kevin Dolan)

TecnoTut,

Quote:

I have never met a scientist who works in the field of consciousness, who thinks that consciousness is causally

inefficacious.

Is there a law in physics that states interactionism is false? No. Is there a law in physics that states minds are merely brain activities? No. Only philosophical theories, not scientific ones, claim interactionism is false and that minds are brain activities.

Well, I certainly don't know of any scientific theories which hold that interactionism is false, but there are most definitely scientific theories holding that the mind is brain activity.

Quote:

That would be insane, since it would imply that consciousness cannot be studied scientifically.

Scientifically, one can study what a bat's brain does, but we still would not know what it is like to be a bat. So in that sense, consciousness does indeed defy scientific inquiry.

There is absolutely no reason that science should be able to cause a human being to know what it is like to be a bat. As I have already pointed out in this thread, knowing what it is like to have an experience, is not an example of knowing abstract facts about the world, which is what science can provide us with. Knowing what it is like to be a bat means both experiencing being a bat, and possessing the memories of having been a bat. Knowing all of the facts about those experiences and memories, is not going to magically cause those brain processes to occur in your brain.

Saying that physicalism requires that science should be able to let me know what it is like to be a bat, is like saying that it requires that a scientific description of a car should cause a real car to appear in my driveway. It is pure nonsense, derived from the conceptual confusion between the thing being described, and its description.

Monroe,

Quote:

How can an experience be physical? It may be caused by a brain process, the outcome of a process, but it cannot be a process. The idea is incoherent. (Monroe)

How so? You keep saying this, but have offered no explanation why. (DM)

How is it coherent? Let's say I just don't understand how what you are saying makes sense (which I don't). How are you going to explain it to me, other than saying, "It just is!"? (Monroe)

What exactly are you having trouble understanding? Is the statement "experiences are something that your brain does" not clear?

Quote:

Remember, the issue here is not whether it is possible for there to be some hypothetical world where brain activity does not cause our behavior, but instead spirits do. The issue here is whether or not our world is such a world. Science indicates that our behavior is caused by brain activity. We can either accept science, and reject the superstitious notion of spirits, or we can reject science, and believe any logically self-consistent hypothesis which makes us feel good. (DM)

Another causal option is that where we think there is a closed causal structure of physical events in the brain, there are actually mental events in between. Instead of $P1 \rightarrow P2 \rightarrow P3$, it is $P1 \rightarrow M1 \rightarrow P2 \rightarrow M2 \rightarrow P3$. (I think you need to be more careful about looking at all the logical possibilities.) It's not that the brain doesn't cause behavior, it's just that something nonphysical is involved in the process. So it isn't that the brain doesn't cause behavior, it does, and consciousness does also since they have causal interactions. (Monroe)

This does not address my point. If science tells us that the chain of causality is $P1 \rightarrow P2 \rightarrow P3$, then of course it is possible that the actual chain of causality is something else. But there is absolutely no reason to posit the existence

of any other causal factors. There is no way to distinguish between any of the infinite number of other possibilities. Either we can accept the answers that science gives us, or we can reject them, and say that the chain of causality is unknowable.

Again, I am not going to reject the answers which science gives me just because they are a bit counter-intuitive.

Quote:

The motivation for this is that if we only talk about physical events, it seems that we are leaving out consciousness.

It may seem that way to you, but I maintain that such a position is incoherent. It is not even *conceivable* that a *complete* description of *all* physical events could *possibly* leave out any of the phenomena which I think of as my consciousness. Any description which left them out could not possibly be complete, because the effects which I know my consciousness has would not be included in such a description. At worst, my complete physical description could be *wrong*. That is always a possibility. As you said, science could tell me that my actions are caused by brain processes, when they are really caused by some spirit. It would then be correct to say that my scientific description of the world is *wrong*, but it would not be correct to say that it leaves out consciousness. It does not leave it out. It merely misinterprets what it is.

I am fully aware of the fact that science will give me wrong answers if the axioms upon which it is based are false. There is nothing I can do about that. There is also nothing "special" about consciousness in this respect. The exact same criticism could be made of *any* scientific theory. Maybe gravity doesn't cause the Earth to orbit the Sun? Maybe Jesus causes it through sheer force of will? But I am not going to toss out all of science just because I cannot logically prove that it cannot possibly be wrong.

Quote:

In order to support the claim that we are not leaving out consciousness, you would have to show that consciousness simply "is" some of those physical things, events, or processes.

I need do no such thing. I need only show, as I have done above, that for a complete physical description of the world to leave out consciousness, would be logically self-contradictory. Again, that description may end up being wrong, but it does not leave out anything which I think of as consciousness, or any aspect of the mind which I could possibly know that I have.

Quote:

This is not trivially true, since the notions of consciousness and brain events have completely different senses. So, if true, it must be that by a careful analysis, you could deduce the identity. It cannot be contingently the case that consciousness is a brain process, because you would be leaving out information about consciousness by only describing brain processes. (Contingent things aren't deducible.) If you want to support the deduction by bringing in empirical correlations between conscious states and brain states, then you are not being reductive about conscious states because, simply, you are bringing in information other than what is obviously physical information to deduce consciousness. (You'd have to assume your conclusion to do that.) And an empirical correlation will not support an identity claim. So it must be a priori and necessary, if true at all, that consciousness is identical to brain processes. But this is obviously not the case, as you have conceded.

Again, it appears you are talking here about metaphysical identity, and metaphysical reductionism. Neither of these things have any place in science. Neither of these things are being, in any way, advocated by me. I am talking about a scientific *description* of those aspects of my mind which I know I have. Since I know that I have them, I know that they have physical effects. This means that as long as the axioms of science are valid, I should be able to construct such a scientific description. Such a scientific description says absolutely nothing about metaphysics or ontology.

As I already explained, it could very well be that minds are the true ontological existent, and that they created the physical universe as a kind of illusion, and that our brains are simply the projections of those minds into that physical illusion. Again, so long as the axioms of science are met, and this illusionary world functions according to natural laws which can be determined from observations made within it, science works just fine. And if it turns out that in this idealistic world, every aspect of our own minds which we are capable of knowing we have, can be explained in terms of these brain projections, then the scientific theory of the mind as brain processes works just fine too. I have absolutely no interest in what the actual metaphysical or ontological state of affairs is. It is unknowable, and completely pointless to worry about it. What I am concerned with is understanding the world around me to the

best of my abilities. Consciousness (at least those aspects of it which I know I have), is clearly a feature of this world, and brain activity seems to be the best available explanation for how it works.
By Death Monkey (Kevin Dolan)