

Consciousness and Neural Activation

James Blackmon
Spring, 2013

"I do not wish to give the impression that I think there is no mystery about consciousness. There is, for instance, something of a paradox connected with any attempt to localise it."

Alan Turing

1.

A common view in the sciences is that the human brain, like any other molecular entity, is a scattered material object. As dense as it is and as continuous as it might appear to be, we know the brain to be largely a tangle of neurons separated from each other by synaptic gaps across which volleys of neurotransmitters are cast back and forth. Moreover, owing to modern physics, we understand that at a much lower level the human brain, as well as each of its neurons, is a constellation of atoms and subatomic particles, each such particle separated by comparatively vast expanses occupied by nothing we would count as material.

Neuroscience tells us that this particular object is the thing that is conscious. As philosophers we say that it instantiates conscious thought, that it or some of its properties realize conscious thought, or that conscious thought supervenes upon some of its properties. Stated in terms of supervenience, we tend to agree that, for each conscious mental property, there are physical properties which, if instantiated by some brain, provide sufficient conditions for the instantiation of that mental property by that

brain.¹ This is a general supervenience thesis upon which a wide variety of theories of mind (physicalism, functionalism, property dualism, emergentism, and epiphenomenalism) can agree.² Furthermore, the general consensus is that the relevant physical properties involve neural activation. So our general supervenience thesis is wedded to the commonly held thesis that the base of supervenience is the activation of neurons in the brain, a scattered material object.

Using a thought experiment, I will argue that this neural activation view leads to some unexpected consequences.³ These are consequences which many people will see as practical (but not logical) absurdities. After developing the consequences of the standard view of neural activation, I will suggest an alternative to this conception of the mind for those who would like to avoid these consequences.

Two qualifications are in order. First, although my admitted inclination is that the consequences are absurd, or at least that they are very strange consequences for a theory of mind, this might be just an inclination; I won't try to convince readers to share it. My primary aim is to reveal some unexpected consequences of the standard view

¹ Because I am interested in mental properties as individuated by their phenomenology, I adopt a narrow content view. This is not a rejection of the view that mental properties may also have broad content.

² If we keep from stipulating that supervenience can only hold between non-identical properties, physicalism is committed, if only trivially, to supervenience. Typically, non-identity is not required in formulations of supervenience.

³ The thought experiment follows a long tradition of ignoring what may appear to be hard technological limits on what humans can do in order to get entities other than brains to mimic brains. Ned Block's "China Brain" [1978], Daniel Dennett's "Hamlet and Yorick" [1981], Arnold Zuboff's "Spread Brain" [1981] and the "Neuron Replacement" thought experiments referenced by Zenon Pylyshyn [1980], John Searle [1992], and David Chalmers [1996] are a few such cases. In doing so, I believe I am in the good company not only of these philosophers but of physicists and mathematicians who also expect their theories to be capable of addressing highly unusual but physically possible situations. Physics and mathematics can address the systems I will be describing; we should expect philosophy of mind to do the same.

that consciousness supervenes on neural activation, and to provide an alternative.

Second, while the extant target of my paper is this standard neural activation view, the argument generalizes to views on which consciousness supervenes on the physical properties of any collection of physically discrete objects: molecules, atoms, and so on. This will be elaborated in the final section of the paper.

2.

Consider some brief but complete conscious mental event in your life, this one you are having now as you read these words complete with all its sensations, thoughts, and emotions. Call this mental event E . According to the standard view, E corresponds to some pattern of neural activation, N , such that any brain having N will instantiate E . Now consider the following question: Is the mereological sum of these neurons exactly that entity which instantiates the mental event in virtue of having the right intrinsic properties?⁴ Suppose this is the case.

Consider now a system of neurons qualitatively identical to the system that composes your brain while instantiating N , but for one neuron, k . The difference is that neuron k sits inches outside the rest of this brain, and yet it is linked to its corresponding neighbors with circuitry that maintains the appropriate signal transfers. That is to say that despite k 's unusual location with respect to the rest of the brain, it behaves

⁴ Putting the question this way presupposes that mental states are instantiated and that instantiation is a relation between the mental state and some thing. Perhaps there are reasons to avoid adopting this ontology. If so, we can drop that presupposition by asking a more general question: Is the mereological sum's having N a sufficient condition for the occurrence of the mental event? But for convenience, I will often speak in terms of instantiation with the understanding that I am not insisting on this ontology. My claims can be recast in the more general sense which appeals only to sufficient conditions.

neurologically just as it would, had it been in its proper spot.⁵ Of course, as part of such a system, neuron k could be even farther away, perhaps miles, as long as the signal transfers were reliable and appropriately timed, and even at this distance, k would behave as it would had it been in its proper spot.⁶

We can now consider an alternate system much like this one but differing in one very significant respect. For, note that the holding of causal relations between neuron k and its neighbors is not a necessary condition for the proper behavior of k . So, instead of one circuit between k and its neighbors transmitting the requisite signals, a more complicated system may exist, one that controls the firings of k and k 's neighbors externally, ensuring that k exhibits the appropriate neural activity. This system would consist of a master computer (MC) with control filaments running from MC to k and k 's neighbors. No signals would be sent between k and the neighbors. There would not even be signals sent between them *via* MC ; MC would not be a "bridge" between neurons, nor would it take any input. Instead, MC would simply be programmed to "puppeteer" the neurons, to send the requisite signals out along its control filaments at

⁵ Of course, creating such a system would be an accomplishment. But the system just described is physically possible, and a complete theory of mind (just as a complete theory of physics) must admit it. The job of physics, in the limit, is to be able to answer all physics questions about all physically describable systems including this one. How much would it weigh? How would it radiate heat? What would happen if it collided with another such system at 40 m/s? Similarly, the job of a theory of mind, in the limit, is at least to be able to answer all questions about the mental properties of all physically describable systems.

⁶ Time-lag due to distance would not be a problem. Neural signals are notoriously slower than electrical signals or those sent along fiber-optic cables. So, a long fiber-optic circuit could match pace with a short neural one. Also, the transportation of actual neurotransmitters is not required. Stores of the requisite molecules can be on hand at each end of the circuit so that when neuron k expels some neurotransmitters that would have gone to a neighbor, a mechanism at the other end of the circuit replicates this by expelling some from its one store. The command to do so would simply be part of the overall signal.

just the right times so that the entire collection, k included, would instantiate some portion of N , a mere pattern of activations which are now occurring as the collateral effects of MC's signal-sending activity. The specific pattern of activation that holds of k and any one of k 's neighbors would be achieved by the *projection* of signals from MC onto each neuron, not by the *exchange* of signals between them. Thus, these neurons interact no more than do the light bulbs of a theatre marquee or the pixels of your computer screen.

There is nothing in physical principle to stop such an arrangement from being applied to all the neurons of a collection. Consider each neuron now causally isolated from its previous neighbors but firing appropriately due to the signals it receives from a filament running to a reprogrammed MC, an MC programmed to direct the firing of every neuron in the entire collection in just the right way.⁷ This is yet a collection having the pattern N . Call this thing a *de-integrated brain*. The neurons of the de-integrated brain are all neurons causally isolated from each other, but they nevertheless fire just as they would have fired had no causally isolating de-integration obtained.

Now we confront a question: Is the de-integrated brain a thing that has E , the experience you had previously while reading the beginning of §2? Put another way, if your brain were currently in the de-integrated state but governed by MC so as to exhibit the same pattern of activation it currently has, would you be having the same phenomenal experiences you are having now?

⁷ Note that no robust processing is going on. The manner in which the signals go out may all have been arranged ahead of time. In fact, the system for managing such signals would not have to be very elaborate at all, having the same computational complexity profile as the cylinder of a music box.

3.

If so, we face a number of counter-intuitive results that will be outlined here. Recall that the view under consideration is that certain patterns of activation holding of collections of neurons are sufficient for the occurrence of certain mental events, and the thing that instantiates the mental event is just the mereological sum of these neurons. Now for the results.

First, this de-integrated brain may span from Mars to Venus. In fact, there is no limit in principle to how far the neurons may be from each other. Physical instantiations of mental events can be radically strewn across space, their parts bearing little or no causal relevance to each other. Moreover, these physical systems may have less integrity than any cloud of gas, their elements straying far and wide, flying past each other at irregular speeds. The entire system of neurons can do whatever we might imagine as long as each fires in just the right way at just the right time. For some people, this is absurdity enough; physical “mind-having” things are, for some reason or other, intact things. They do not mix or dissipate. But if this is a bad consequence of the view under consideration, it gets much worse.

For the second result, suppose MC has one extra neuron on hand, and suppose that this neuron is firing in just the same way that your neuron k is firing. We have a new collection of neurons here: your collection, minus k , plus this additional foreign neuron. The collection that is your de-integrated brain drastically overlaps this other collection. In fact, the two collections share all but one of each other’s neurons. Yet, because the two collection-specific neurons behave the same way, both collections

exhibit the activation pattern N . Thus, according to the view under consideration, both of these radically overlapping yet different collections of neurons have the same experience you just had while reading §2. Of course, MC could have many foreign neurons doing just what k or other neurons in your collection of neurons do, and thus there could be many other de-integrated brains overlapping yours, all having the experience E .

Third, that extra neuron might have come from another de-integrated brain, a collection also controlled by MC or even some other master computer. If so, then realizations of mental states just like yours may overlap more than one de-integrated brain. Again, the overlap may involve more than one neuron. Part of one de-integrated brain might be collaborating with parts of other brains in the instantiation of “someone else’s” mental event whether it is one that corresponds to N or not.⁸

Fourth, we must eventually admit that MC along with its control filaments is merely a conceptual aid, one that assisted us in imagining different collections of neurons. Even without it, many collections of neurons exist. So, for instance, unless your neuron k is doing something especially rare, then for some period of time there exist counterpart neurons in the heads of other people (and other animals) and each of these neurons helps to compose a new collection of neurons firing in just the way that your collection of neurons is firing. We do not need MC or cleverly-constructed systems of circuitry in order to establish the existence of collections of neurons that meet the

⁸ I do not mean to unjustifiably assume that every collection has its own subject. Thus the scare quotes.

requisite conditions. The consequence is a bizarre (though limited) form of panpsychism: Parts of our brains collaborate in instantiating many minds and many of these minds have overlapping physical instantiations.

Fifth, and finally, we must ask whether only neurons may form these collections. On many views, the replacement of one of your neurons with something else that behaves in the same relevant way does not disqualify the resulting system from instantiating the mental state. Indeed, on some views, behaving in the same relevant way as the neuron is plainly sufficient for partaking in the instantiation of the mental state.⁹ If so, then our panpsychism runs further amok to involve collections that include other things that act sufficiently like a neuron. Build a single artificial neuron that fires in the right way, and you have created multitudes of overlapping physical things that have conscious mental events.

We have reached the end of this branch. If collections of neurons instantiate mental events merely in virtue of having the right patterns of activations, then some form of panpsychism follows.¹⁰

⁹ The neuron replacement thought experiment [see Pylyshyn 1980] is used to make or test this point. On the views referenced, should a neuron be seamlessly replaced with an artificial functional equivalent, there is no difference in the conscious mental properties of the system, which is no longer a fully organic brain. And should replacement continue until we have nothing but a functionally equivalent artificial system, there is no difference in conscious mental properties. We will return to neuron replacement in the final section.

¹⁰ Absurd as these consequences may be, I do not see this as a *reductio ad absurdum* for three reasons. First, we obviously have not encountered any specific contradiction, just a view that seems very strange, even absurd. Second, surely some truths are very strange, and if any of the previously covered consequences did hold, we could never know it through any scientific means. Do these disparate collections of neurons and neuron-like things have full-fledged experiences? This is just the problem of other minds applied to very unfamiliar things, and I admit that for the true skeptic, the problem of other minds holds for these things just as it holds for anything else. Third, absurdity alone (without contradiction) is no license for rejection. For, the remaining options may be just as absurd or even more

4.

On the other branch of the dilemma, collections of neurons do not instantiate mental events merely in virtue of having the right patterns of activations. One criticism of the thought experiment as it has been developed so far might be that the appropriate kind of neural activation involves more than the right neurons having the right activations at the right times; the activation must *spread* causally among neurons. According to this objection, the brain certainly is not like a mere theater marquee; appropriate activation necessarily involves real signal transfer between units. If so, then I will seem to have been unfair to those who had this kind of phenomenon in mind. For de-integrated brains do not after all have *N*, if *N* is to be understood in a way that necessitates appropriate signal transfer between units.

I will use the term *spread of activation* in the way that involves necessary signal transfer between units. The term *pattern of activation* shall refer (as I have been using it) only to the activation states of the neurons over time, remaining neutral regarding the cause. Because having spread of activation is only one way to exhibit a pattern of activation, I will use *mere pattern of activation* to identify the patterns that do not hold due to spread. Therefore, de-integrated brains exhibit *patterns*, but not *spreads*, of activation among their neurons, and their patterns are *mere patterns*.¹¹ The proposal then is to require activation spread, thereby avoiding all the presumably embarrassing

absurd. If all other options turn out to be significantly more absurd, we may have to embrace these absurdities as the most plausible consequences of our standard view.

¹¹ The fact that signals are transferred from MC to the neurons in the relevant cases should not be misconstrued as satisfying the condition that signal transfer must hold between the neurons.

panpsychic results just witnessed. Certainly, the previous thought experiment does not involve collections exhibiting spreads of activation; we have only mere patterns. But does anything other than an *ad hoc* attempt to avoid these consequences support this proposal? Let's look more closely at the notion of activation spread.

Either some appropriate scattered mereological sum is the thing that instantiates the conscious mental event or not. If it is, then whether it instantiates the conscious mental event depends not only on the pattern of activation it has but on whether its neurons exchange signals in the appropriate way. But whether this condition holds will have very little or no physical effect on any of the neurons. Physical changes in any particular neuron are brought about by the immediate event of some signal's reaching that neuron; they are not brought about by the historical event of that signal's having been sent from a legitimate (i.e. proximal, incorporated, and sufficiently neural) source. So, on this view, whether any neuron plays any role (typically a collaborative role along with other neurons) in the instantiation of some conscious mental property depends on features of the world that have little or no physical effect on any of the neurons.

One might object that nevertheless when the signal arrives, the identity of its source has had some appreciable effect on the signal, an effect that is in turn the cause of an appreciable effect on the neuron. That is to say, different sources cannot send sufficiently similar signals nor can they bring about sufficiently similar effects in neurons. Perhaps the argument can be made that the neurons of intact brains and de-integrated brains are, for physical reasons, necessarily different at the thermodynamic level and

that these differences hold despite similarities in firing patterns. If so, the relevance of such differences still needs to be established, hopefully in a manner that avoids being *ad hoc*.¹²

But if relevant differences in signals and neurons do not exist at any sub-neural level and yet signal source matters to whether the mental event is instantiated, then we have a view of consciousness that looks like magic. The instantiation of the conscious mental event *E* is determined by factors by which the thing doing the instantiating is not physically affected. After all, the de-integrated brain and the intact brain are neuron-for-neuron qualitatively identical; no neuron behaves any differently as a result of its degree of spatial or causal intimacy with another. On this view, conscious mental states do not supervene simply on the intrinsic properties of the material thing, a scattered object, and thus, consciousness does not locally supervene. Non-intrinsic (albeit intuitively “internal”) features of the scattered object matter, even when they leave no physical trace on the scattered object that instantiates the mental state.¹³

¹² As I understand it, this is a nontrivial project, and the faith we place in it will draw heavily from our knowledge of thermodynamics and signal theory. For whether and to what degree there are bound to be differences between the signals from these two different sources (differences that do not affect activation) is an issue of physics. This is not to count the problem as a one of pure physics. After all, whether and to what extent any such differences are relevant to the instantiation of a mental state cannot be solved with equations. But I think it is safe to say that this is unexplored territory. Moreover, it would certainly be a drastic shift in cognitive neuroscience to hold that mental properties are not instantiated at the level of neurons and neural activation but at some deeper thermodynamic level, especially when the relevant deep-level differences *do not have emergent physical effects* on neural activation.

¹³ There is an exegetical problem here. To characterize these non-intrinsic features of the scattered object as “external” or “non-local” threatens to create confusion. For, some of these features can be found interconnecting the parts like bridges interconnecting islands, not strictly existing within any part (island) of the scattered object (the sum of the islands), but within a continuous region that circumscribes it. As such, the features are intuitively “internal” and “local” according to some ways of thinking about it. My

This would be action at a distance of a special kind. Non-local physical events would make a significant phenomenological difference to a thing that was not physically affected by those non-local physical events. Following Jaegwon Kim, let's admit that Xantippe becomes a widow the moment Socrates dies.¹⁴ She gains the relational property of widowhood, a property which supervenes, not locally, but globally, depending on a property (*being dead*) of something external to her (Socrates). Indeed, this property will become relevant to future causal processes; however, let Xantippe and Socrates be sufficiently far apart at the moment of his death (on different continents or planets) and we can see that this property is causally irrelevant to Xantippe at the moment of its instantiation. Xantippe gains the property of widowhood, but she undergoes no physical change that would not have occurred had she not gained that property.¹⁵

Science cannot hope to study such causally irrelevant relational properties merely by studying the features or behavior of the things that have them. This is simply because having causally irrelevant relational properties is irrelevant to a thing's physical features or present behavior. Thus, if signal source matters despite its physical irrelevance, then having *E* is not like widowhood, for while widowhood itself brings

point is that such features are nevertheless not relevant to the intrinsic physical nature of the de-integrated brain.

¹⁴ Cf. Kim [1993].

¹⁵ Moreover, there is no sense in which even an omniscient being can discern in just her physical make-up some Leibnizian mark of widowhood.

about no detectable change in the person, having *E* is supposed to bring about a very significant change, a change in phenomenal consciousness itself.¹⁶

For illustration, consider one last extension of the thought experiment. Suppose that each of your neurons is causally isolated and governed by MC as before but the neurons are not moved away from each other in any significant way. On the current view in which spread of activation is necessary, you do not have *E*, for there is no spread of activation. Suppose, however, that MC removes these causal barriers and simultaneously cuts its signals. We are supposed to accept that you instantly become phenomenally conscious as a result, “waking up” in state *E*, effectively going from philosophical zombie to human in an instant. And this very significant event occurs even though not a single neuron has diverged in either its state or behavior.¹⁷

So far, we have focused on the view that conscious mental properties supervene on the neural activation of the brain, a scattered object. The de-integrated brain thought experiment requires us to distinguish between a mere pattern of activation and a spread of activation. If we only require a mere pattern of activation, de-integrated brains still give rise to conscious mental properties. Yet, additional consequences follow:

¹⁶ Putting it in other terms, consciousness on this view is the product of a mere “Cambridge property” as recognized by Peter Geach. See Shoemaker [1980].

¹⁷ Moreover, once we permit a material thing’s having consciousness to depend on its causally irrelevant relations to other things, it is not clear that we can avoid competing potential bases of supervenience. Why hold that it is the causally unaffected brain that is the base? For example, one might suggest that instead there is a special neuron, the one and only neuron which instantiates your conscious mind—so long as it is appropriately related to all other neurons regardless of whether these relations have any physical effect on the special neuron. Or one might suggest that there is a special grain of sand, the one and only grain of sand which instantiates your conscious mind, so long as it is appropriately related to all the relevant neurons regardless of whether these relations have any physical effect on the special grain of sand. In sum, this kind of dependence on causally irrelevant features opens the floodgates for uncounted futile proposals.

the “brains” which give rise to conscious minds are widely scattered, grossly overlapping, and causally disjointed. Any view according to which a mere pattern of activation is sufficient for consciousness must confront these consequences. If on the other hand, spread of activation is required—and yet the brain continues to be understood as a scattered material object, then consciousness depends on strictly external factors which are (thermodynamic considerations aside) physically irrelevant to the physical features of the instating object.¹⁸ The next and final section proposes and briefly considers a remaining option: that the supervenience base of consciousness is a non-scattered object.¹⁹

5.

Perhaps we can avoid the allegedly absurd consequences by rejecting the thesis that it is just the scattered object that instantiates the conscious mental state. Preferably, we’ll be able to do this without stipulating any nonphysical entities. One suggestion is that mental states are instantiated by the sum or composition of both the relevant material and whatever it is that keeps it appropriately connected and in causal contiguity—that medium or field by which the appropriate activations actually spread. Instantiations of minds, on the view suggested, are physical in the sense that light,

¹⁸ As noted in footnote 13, these “external” factors hold between the parts of the scattered object. They are external to the parts, but internal to any of the continuous regions we intuitively tend to recognize as including the scattered object.

¹⁹ Being an investigation of the consequences for supervenience (on the physical), this paper leaves aside the option that there is no physical supervenience base. Similarly, I leave aside the option that while there is conscious thought, no physical entity instantiates it. This option is taken in idealism and solipsism.

gravity, and space-time is physical. They just are not entirely material in any classical or intuitive sense.²⁰

Of course, whether it makes sense to speak of some medium or field existing throughout what has often been thought of as empty space is largely a matter of physics. But however it is that we can best characterize that which mediates the transfer of energy or the causal connection between material things, this final attempt to avoid the strange consequences just surveyed appears to have some promise. For, if it makes sense to say that the signals themselves, along with the material, are collectively the bearers of mental properties, we can avoid all so far witnessed consequences. After all, on this view, the thing that has conscious mental properties has those properties in a way that does not depend on external facts, not even the “intuitively internal” external facts holding of the spaces between parts of a scattered object. Consciousness, on this view, is perfectly intrinsic to the entity that has it; there are no qualitatively identical things bearing different physical relations which differ regarding conscious mental properties, no things whose consciousness can be instantaneously activated or deactivated without being physically affected or changed intrinsically. So, to that extent, there is no phenomenally effective Cambridge change, no “magic” as previously construed.

²⁰ Nor are they entirely material according to a prevalent, surviving view in contemporary science. After all, some parts of the things which instantiate minds, on this view under consideration, fill the “empty” regions between and neurons, and these regions are not composed of atoms, nor are they commonly believed to have any mass.

Moreover, if it makes sense to say that the signals themselves—and nothing else—are the bearers of conscious mental properties, we also avoid all so far witnessed absurdities.²¹ I suggest that functionalism, at least at the neural level, encourages this view. On such a view, other features of the material are irrelevant so long as the material can perform its proper function, and this function is simply the adequate transmission of energy signals. The neuron replacement thought experiment, beginning with Zenon Pylyshyn [1980], provides conceptual support for functionalism at the neural level.²² According to the thought experiment, neurons are replaced by artificial functional equivalents.²³ The point is to show that it makes no difference to the surrounding neurons whether they exchange signals with a neuron or an artificial replacement (or anything else) as long as the replacement is functionally equivalent to the replaced neuron. If conscious mental properties are realized by the physically relevant functional activity of the neurons, then despite these neuron replacements, conscious mental properties persist just as they would have had no such replacements occurred. On this view, what matters is *how* things are causally affected, not *what kinds of* things are causally affected. But how things are causally affected depends on the

²¹ Perhaps the view under consideration accords with the view that consciousness is an electromagnetic phenomenon, as advanced by Pockett [2000] and McFadden [2002]. McFadden, however, claims to be addressing the Hard Problem of consciousness. I am making no such claims here, and I don't see how the Hard Problem can be solved simply by identifying the supervenience base with an EM field, or with anything, material or otherwise.

²² See also Savitt [1982], Cuda [1985], Searle [1992], and Chalmers [1996].

²³ If one balks at the idea of physically replacing a neuron without disrupting the system, we can avoid supposing that any disruptive replacement occurs, stipulating a series of distinct "brains" each one more artificial than the previous one, but otherwise having any remaining neural portion bear qualitative identity relations to the previous "brain".

signals those things receive. Dependency then reduces to whatever it is that signals are or to whatever it is that mediates them, again an issue that is largely a matter of physics.

The point here is not to determine what signals are or what energy is in some physical or metaphysical sense, but to show that on the view under consideration there are roles to be played, one that is often characterized as “signaling”, “the transfer of energy”, and so on.^{24,25} In physics, the role is the bringing about of innate changes in neighboring, if not contiguous, entities, and perhaps *being* those changes, as well. In our theory of conscious thought—one respecting supervenience but seeking to avoid the previously witnessed results—the role is instantiating or providing the supervenience base of conscious mental properties. Thus, unless we are going to cling to the “neuro-chauvinist” dogma that sufficiently neural material is necessary for conscious thought, a commitment to the idea that physical entities instantiate conscious thought and a commitment to avoiding the alleged absurdities witnessed above together entail a positive thesis: Conscious thought is, or is instantiated by, continuous regions of energy, appropriately channeled.

²⁴ Energy, as typically defined in science textbooks as the *capacity* to do work, is strictly a dispositional property. Of course we might also interpret attributions of energy to refer not to the dispositional property but to something in a system which grants that dispositional property. Or we may try to see it only as a quantitative measure, one that has complex and infinite conditional implications. Surprisingly, while related notions such as causality, counterfactuals, matter, and dispositional properties have received much attention in metaphysics and the philosophy of science, few philosophers seem to have much to say about energy.

²⁵ As Skyrms [2010] shows, very much can be discovered about signals without trying to say much about what they are or what may bear them. And these further questions can be avoided while nevertheless recognizing the fundamental role of signals in thought: “Signals operate in networks of senders and receivers at all levels of life. Information is transmitted, but it is also processed in various ways. Among other things, that is how we *think*—just signals running around a very complicated signaling network.”

6.

How strange is this view? Possibly not as strange as it may sound at first. It may sound dualist or “spiritual” or just flaky to say that conscious thought is energy, but we should remind ourselves how little this says given our current understanding of energy, something on which I am taking a lot of pains to remain neutral. The positive thesis I am considering here is simply that the base or instantiation of conscious thought is continuous, and in portions not material—if the term ‘material’ retains any restrictive application these days.

Currently, our best scientific theories discourage us from maintaining a principled distinction between mass and energy. If this is right, then the implications I’ve drawn out here may not weigh heavily against so-called materialist theories of the mind; however, this would only be because there would be no coherent strictly materialist view of the mind to hold.²⁶ If so, then we are left with two basic views regarding the supervenience base of conscious thought. One view, what I’ve referred to as the standard neural activation view, sees the physical entity that instantiates consciousness as whatever it is that fills all and only the scattered region we typically associate with the mereological sum of a brain’s neurons. The other view sees the relevant physical entity as filling a region that is continuous. We may think of the region as filled with some material along with something else, but I don’t see that we have to insist on any

²⁶ Contemporary definitions of ‘materialism’ no longer have much to do with any interesting conception of matter, for the definitions defer to a physics which includes numerous entities that are not considered to be material and that are believed to lack the typical characteristics (mass, extension) of matter. For this reason, ‘physicalism’ is a more accurate term.

particular distinction between material and immaterial things, and the issue ultimately appears to be independent of whether the entity that instantiates conscious thought is scattered and whether this depends on physically irrelevant external conditions as investigated in §4. I have argued against the view that mere patterns of activation are sufficient for consciousness, positing a continuous entity as the supervenience base of conscious thought.

Acknowledgment: Thanks to Jonathan Dorsey, Robert Howell, Michael Jubien, and Arnold Zuboff, and for helpful comments on earlier drafts of this paper.

Works Cited

- Block, N. (1978), Troubles with functionalism, in C. W. Savage (ed.), *Perception and Cognition: Issues in the Foundation of Psychology*, Minneapolis: University of Minnesota Press.
- Chalmers, D. (1996), *The Conscious Mind*, New York: Oxford University Press.
- Cuda, T. (1985), Against neural chauvinism, *Philosophical Studies* 48:111-27.
- Dennett, D. (1981), *Brainstorms*. Cambridge, Mass.: MIT Press.
- Kim, J. (1993), *Supervenience and Mind: Selected Essays*, Cambridge University Press.
- McFadden, J. (2002), The Conscious Electromagnetic Information (Cemi) Theory Field Theory: The Hard Problem Made Easy?, *Journal of Consciousness Studies* 9, No. 8, 2002, pp. 45-60.
- Pockett, S. (2000), *The Nature of Consciousness: A Hypothesis*, Lincoln, NE: Writers Club Press.
- Pylyshyn, Z. (1980), The "causal power" of machines, *Behavioral and Brain Sciences* 3:442-44.
- Savitt, S. (1982), Searle's demon and the brain simulator reply, *Behavioral and Brain Sciences* 5:342-43.
- Searle, J. (1992), *The Rediscovery of the Mind*, Cambridge, Mass.: MIT Press.
- Skyrms, B. (2010), *Signals: Evolution, Learning, & Information*, Oxford.
- Shoemaker, S. (1980) Causality and Properties, *Time and Cause*, Peter van Inwagen (Ed.), Dordrecht: Reidel.
- Turing, A. (1950) Computing Machinery and Intelligence, *Mind*.
- Zuboff, A. (1981), Hofstadter, R. and Dennett D. (eds.), *The Mind's I: Fantasies and Reflections on Self and Soul*. New York: Bantam Books.