

Discussion on Dennett, Philosophy Forum, Apr 2005

Subject: Interview with Dennett

The following is a quote of Dan Dennett from an interview with Robert Wright. The entire interview is in video format and can be accessed through the following website: <http://meaningoflife.tv>

"...what consciousness IS is the relative political influence or fame of structures in the brain that win out in competition against rival structures for domination of the brains activities in various ways - thats putting it very programatically, but basically it is saying that in your head there's a sort of turmoil going on, ...and there's many different contentful events vying for king of the mountain, vying for control. And the ones that win - by default something always wins when you're awake - that's what YOU'RE CONSCIOUS OF"...And it's not that when it wins THEN consciousness happens or that it kindles some further thing that is conscious; THAT'S JUST WHAT CONSCIOUSNESS IS"

-Dan Dennett

I find the above quote a little misleading. First he seems to be saying that the structure of the brain that wins in competition IS what consciousness is, but then he says 'that's what you're conscious OF' implying that the structure IS the CONTENT of your consciousness. But then he goes back to saying 'that's just what consciousness IS'. Is there anyone out there who wants to clarify what Dennett is trying to say in the above quote?

By Decartes1

I think that one of the things he is trying to convey is that this somewhat dualistic notion of "structure of consciousness vs content of consciousness" is misguided, and does not really have any relevance to what is actually going on.

Notice also that he says that it is the *events* which are vying for control, and that it is the winning *event* that "is" consciousness at that moment. Not the brain structure. These events are processes being performed by the brain structures.

By Death Monkey (Kevin Dolan)

Descartes1 wrote: But it seems to me that Dennett believes the events in question are structures of the brain and not a product of these; but again, I'm not so sure.

Hmmm... I certainly don't get that impression from what he said. What specifically gives you that impression?

By Death Monkey (Kevin Dolan)

Descartes1 wrote:

From the first sentence in the quote. He says 'what consciousness IS is the relative political influence or fame of structures in the brain that win out in competition against rival structures.

Exactly. He says that it is the relative *influence* of structures in the brain (which refers to the processes those structures are performing), not the structures themselves.

By Death Monkey (Kevin Dolan)

What I find especially amusing is that Wright went on Beliefnet after taping this and declared that he had gotten Dennett to declare that “life on earth shows signs of having a higher purpose.” This bit of distortion was then seized upon by conservative religious bloggers as evidence of yet another atheist turning into a believer. Dennett was quick to point out this was completely bogus, of course.

By Faustus (Brian Peterson)

Descartes1,

DM wrote: Exactly. He says that it is the relative influence of structures in the brain (which refers to the processes those structures are performing), not the structures themselves.

Descartes1 wrote: Ok, if you don't mind, elaborate. Give me an example of a process these structures would perform.

I thought that would be obvious: Neuronal activity.

Or you asking for specific examples of processes which certain brain structures perform? If so, I don't think we really have specific names for them. We just refer to them as the brain activity being performed by a particular brain structure.

It seems to me that what Dennett is saying is that various parts of the brain are competing for domination of the brain's overall activity, and that at any moment in time, your consciousness is essentially the process being performed by the part of the brain which won. Of course, when you consider the degree of connectivity between the various parts of the brain involved, it is not an all or nothing situation. You consciousness will be primarily whatever process is dominant, with influences from the other processes which are not dominant.

By Death Monkey (Kevin Dolan)

Descartes1 wrote:

By the way Faustus, do you care to explain Dennett's quote for me? I noticed you were the author of the chapters summaries of 'Consciousness Explained'.

Well, first off, please give some credit to Probeman's entries! We're taking turns at this point; I'm now struggling with Chapter 10, which I personally find the weakest, hence my lack of enthusiasm.

Anyway, my interpretation doesn't differ much from Death Monkey's. The discussion is sort of nit-picky in my opinion. There are a lot of brain processes that operate on a winner take all system, so for instance in experiments on binocular rivalry (two different images, one per eye), you will be conscious only of one image at a time, as the networks representing it “win” or gain more “fame” over their rival. But even in these cases, processes that don't win will still leave a causal trace, so it really isn't all or nothing as far as keeping all the subsequent effects in mind. You might not be conscious of something, but you may nevertheless perform better than chance in a forced guess if it was processed to some degree. It's this further effect of processing that causes Dennett to use such terminology as “relative influence”.

By Faustus (Brian Peterson)

Monroe wrote:

How does an event "win"???

I often hear the process referred to as a “Darwinian” struggle between networks of neurons. This is from different sources, not just Dennett. Apparently it begins in the womb—the brain grows too many neurons, and they immediately start competing with each other for space and activity (neurons that don’t get a constant influx of signals eventually die).

Anyway, it ultimately means dominance over the brain by a pattern of information that begins in one network and spreads, in the sense that the content replicates itself and/or is globally accessible. But ultimately, if that content doesn’t make it to short term memory and the motor areas for a significant length of time, then the subject can’t report it, and it doesn’t become a feature of consciousness.

Descartes1 wrote:

So then if, at any given moment, consciousness IS the relatively dominant process(es) (neuron activity) performed by particular structures of the brain (neurons), then anything, say a robot, that instantiates a functional system like these structures of the brain will be conscious. Would you say this is a correct inference from Dennett's view?

Doesn't he say as much during the interview? At least, I have a possibly hallucinated memory to that effect!

Descartes1 wrote:

Also, how is it that a dominant process amounts to a conscious state?

Well, you can hardly be conscious of something which never makes it into memory and which no part of your body can respond to, so by coming to influence memory and motor areas to a sufficient degree, that’s how content becomes conscious. (This isn’t just Dennett’s approach, but is a hallmark of the global neuronal workspace model, of which his is just one version.)

It’s important to keep in mind that since he’s a type of functionalist, the only sort of theory you’ll get from him is one in which a state’s being conscious means that it has some sort of causal effect. He doesn’t have sympathy for consciousness being defined any other way, hence the attack on epiphenomenalism in the interview. (In other words, once you’ve described all the causal effects a state’s being conscious can have, you’ve described everything it is. At least, that’s how he sees it. Fans of Zombies see things differently.)

Descartes wrote:

My trouble is still with Dennett saying that a dominant event is what you are conscious of but then saying "that is just what consciousness is." Was he just being sloppy with his wording?

No, he really means it!

He goes over this issue and further details of the workspace model in the concluding essay of *The Cognitive Neuroscience of Consciousness*, edited by Stanislas Dehaene as a special issue of the journal *Cognition*. The article itself works better in context, but luckily he’s made it available online:

<http://ase.tufts.edu/cogstud/papers/cognition.fin.htm>

By Faustus (Brian Peterson)

Death Monkey wrote:

It seems to me that what Dennett is saying is that various parts of the brain are competing for domination of the brain's overall activity, and that at any moment in time, your consciousness is essentially the process being performed by the part of the brain which won. Of course, when you consider the degree of connectivity between the various parts of the brain involved, it is not an all or nothing situation. Your consciousness will be primarily whatever process is dominant, with influences from the other processes which are not dominant.

Without having read the whole thread (Faustus kindly just sent me a notice of the discussion) I just want to jump right in and point out that one of the interesting implications of what DM describes above is that much if not all of the processes which form our "consciousness" are below the level of "awareness" and that indeed there is likely to be two way interaction between dominant and less-dominant processes.

In a similar line of thinking I also wanted to mention that in a personal communication Dennett said he "agreed mostly" with the idea that these competing processes in the brain can be seen as a sort of natural selective interaction between memes and sub-memes. I will add that in a sense these meme processes in our brain are analogous to the mostly native but sometimes exotic bacterial fauna in our gut that also compete with each other in a competitive and yet cooperative manner for mutual benefit.

By Probeman (John Donovan)

Muxol wrote:

Exactly how are the influences of the various structures of the brain competing and how does one's influence win over another's (and what are the structures influencing -- each other)?

The descriptions I've heard—which incidentally do not come from Dennett, but from scientists—are that networks of neurons will begin to synchronize their patterns of firing, with the different patterns representing different sorts of content. Competition occurs in the sense that each pattern replicates itself, but in order to do so, it needs to recruit more and more neurons. Alas, there are only so many neurons available, some of which may be part of rival coalitions. That's why this is often referred to in Darwinian terms.

If I recall correctly, in one paper William Calvin uses the example of an ambiguous stimulus which the brain is trying to resolve. After the most simple aspects have been processed, different networks in the higher visual centers begin to go into states of inconsistent content discrimination—"apple" vs. "orange" vs. "ball" or something like that. If one of those contents out-replicates the others and achieves global access for a significant amount of time, that content will be what you experience to the exclusion of the others. If none of them ever gain dominance, you will not be able to report a definitive judgment, but these contents would occur to you as good guesses and the appearance of the object would seem like some sort of morphing of all three.

In most cases, this happens so quickly that ambiguities are resolved before you even knew they existed. Part of the fun of Dennett's *Consciousness Explained* is the way he goes over the experimental literature where the brain is caught in the middle of this process or can be tricked into coming to illusory judgments.

Muxol wrote:

And even if they were competing, and that's all consciousness is, then all consciousness is brain activity. That idea is nothing new.

THAT idea is certainly nothing new. It's been around for centuries. But it's one thing to conclude that consciousness is brain activity, and quite another to have a more specific mechanism to propose. I don't think the idea of computations by Darwinian-like competitions started becoming popular until some time in the late 1980's. Dennett's idea is founded on the emergence of this approach within science, and not from

philosophy.

Muxol wrote:

This idea of competing makes little if any sense.

If the idea made no sense, I doubt very much it would be taken as seriously as it is. Time will tell. The idea in one form is already being put to use in industry, where genetic algorithms for different designs are set against each other in Darwinian competitions, producing novel and efficient engineering tricks the programmers themselves did not anticipate. Evolutionary processes are pretty powerful tools—after all, they made us. No surprise that the brain might employ them, too.

By Faustus (Brian Peterson)

Descartes1 wrote:

Yes he does, but he says that the *robot* will be conscious if it instantiates structures with competing processes like the brain does. But if he believes it is the *processes* that are conscious, then why not say that if a robot instantiates these processes, then the *processes* will be conscious. If he wants to say the *robot* will be conscious, logically, he must be equating the robot to the processes. Do you follow?

To give him the benefit of the doubt, he might say that using the phrase "the robot is conscious" is simply another way of saying that there exists conscious processes within the robot. Now if there also exists, within this same robot, the (illusory) inference that there is a "self" called "robot", the robot might output something to the effect of "I am a robot". And following convention, we might say that "the robot" is conscious. This illusory inference is what Dennett would call the "centre of narrative gravity" - see <http://ase.tufts.edu/cogstud/papers/selfctr.htm>

Descartes1 wrote:

But before I give my reasons, I need to know whether an epiphenomenon is defined such that it cannot, in turn, influence *anything at all*, or just that it cannot influence that which caused it? If the former, then Dennett's criticism seems fine. But if the latter, then Dennett is off the mark. This is because, although a process of the brain produced consciousness, the influence it has on the brain to evoke conversations or form beliefs about it should be an influence on *other* processes than the one that caused it to arise.

Irrespective of how you choose to define the term "epiphenomenalism", you must come to terms with the fact that if an epiphenomenon is what you are invoking as an explanation of the "qualia", then if it has *any* causal efficacy at all, whether it be on the brain process that "caused" it, or on another part of the brain or other processes, then this epiphenomenon is either:

1) a physical process

or

2) something else, which mysteriously interfaces with physical processes

If you do not want to accept 2), then you must accept 1), in which case, the usefulness of the term "epiphenomenon" has been completely dissolved, and we seem forced to talk of consciousness *as being* physical processes.

The only way out of it is to accept that an epiphenomenon is something that has no causal efficacy whatsoever. It is utterly sterile.

But then, as you point out, we are led into the problems of how we can detect qualia.

By [Xeno]Julios

Monroe wrote:

A vying-for-influence seems rather implausible because the standard by which a brain pattern is "fittest" by this account is simply whatever it takes for it to survive. The standard should be that it survives (recurs) because its content is more rational or more useful to the whole organism. This does not imply that various patterns "compete against each other" in any way. Also, the whole idea of competing against each other for influence brings way to much personification/organismification to neural firings.

There is a decent amount of literature devoted to these sorts of issues in the field of cognitive science and artificial intelligence. For a good discussion on connectionism, see Rumelhart (1989) and Smolensky (1989). For an excellent introduction into dynamical cognition, see van Gelder (1996).

Your suggestion that neural competition brings way too much personification into neural firing implies that this competition is directly isomorphic to our rational deliberations. It need not be so.

It might help to keep in mind that there are perhaps many layers of competition, the pruning and strengthening of individual synapses being on a low level, and the relative stability of various larger patterns of neural activity being on a higher level.

Thinking of it in terms of point attractors might also help, where the network is liable to end up in a certain state of activity, given the current state of activation and the current synaptic architecture.

If, on the other hand, your question is how does content, *as qualia* arise out of brain processes, then this is a different issue altogether.

There are many ways in which content can be hypothesized to arise on a *representational* level in brain processing, all the way from symbolic representation, to the weight space in a connectionist paradigm, to the state variables or parameters, or point attractors, in a dynamical system.

Rumelhart, DE (1989). The architecture of mind: A connectionist approach. In MI Posner (Ed.), *Foundations of Cognitive Science*. 133-160.

Smolensky, P. (1989). Connectionist modeling: Neural computation/mental connections. In L. Nadel (Ed.), P. Culicover, L. A. Cooper, R. M. Harnish (Assoc. Eds.), *Neural connections, mental computation*. Cambridge, MA: MIT Press/Bradford. 49-67.

van Gelder, T. (1997). *Dynamics and cognition*. In *Mind design II: Philosophy, psychology, artificial intelligence*, revised and enlarged ed. Cambridge, MA: MIT Press.

All these papers can be found in:

John Haugeland, ed. *Mind Design II : Philosophy, Psychology, and Artificial Intelligence*. Cambridge, MA: MIT Press, revised edition, 1997.

By [Xeno]Julios

Descartes I wrote:

See, this is the problem. You (and Dennett) want to be able to say something like "the processes are what you are conscious *of*" but will also want to say that the processes themselves *are* consciousness. The former way of putting it seems to suggest consciousness is a separate thing of which these processes provide its content, but the latter equates the processes to consciousness. And that was what I was trying to point out earlier when I quoted Dennett and asked you if he is just being careless in his choice of words.

No. I think Dennett is strongly opposed to the idea that "the processes are what you are conscious *of*" because this presupposes a vantage point from which one can view the processes creating consciousness. This is pretty much accepted as an incoherent idea by many philosophers and most cognitive scientists.

Instead what Dennett and others would say is that these processes themselves are what makes you conscious. Which includes perception, memory, talking to one's self and interactive language. Basically they would say that consciousness is the brain's ability to "talk" about perception and memory. The implication of this is that the brain does not have any infallible knowledge of what is going on inside our heads because we aren't conscious of the processes. We are the processes.

As analogy, consider that in a similar way the sensory neuronal processes involved in vision are themselves the color red that we experience or remember. There is no place in the brain that is red just as there is no place in the brain where the Cartesian Theater views these processes creating consciousness. That's an infinite regress.

By Probeman (John Donovan)

Monroe wrote:

Is there any explanation of how patterns have content, or is that beyond the scope of this inquiry?

No, it's a good question but I think you can answer it yourself if you think about why it might be that a ripe berry stands out in a field of green leaves.

Monroe wrote:

What I'm saying is that we may see the effect of a certain pattern recurring more (becoming dominant), but we don't know what causes that. A vying-for-influence seems rather implausible because the standard by which a brain pattern is "fittest" by this account is simply whatever it takes for it to survive. The standard should be that it survives (recurs) because its content is more rational or more useful to the whole organism. This does not imply that various patterns "compete against each other" in any way. Also, the whole idea of competing against each other for influence brings way to much personification/organismification to neural firings.

The cause is perception and memory interacting at the neuronal level to produce behavior that has been naturally selected for survival. After all, most of our behavior can be shown to be influenced by our genetics. But I think all non-scientists would benefit from considering how these processes could create not just stimulus response but complex behavior at a much simpler level. Consider the flat worm brain. It has only some few hundred cells yet can be capable of some relatively complex behavior.

"Planaria (Dugesia, free-living flatworms) are ideal experimental subjects because they are easily cultured and are the simplest organisms with a true nervous system. Planaria are highly cephalized, with a brain to body mass ratio close to that of a rat. They are the lowest organisms with a true synapse and are able to form associations such as between a flash of light and a shock or learn simple tasks such as running a maze. Planaria also have advanced sensory systems, capable of detecting light, chemical gradients, vibrations, electric fields, magnetic fields, and gamma radiation [8].

Planaria reproduce asexually by spontaneous fission and possess an extensive regenerative capability. If a planarian were cut into as many as fifty pieces, each piece could reconstitute its lost

body parts. This property allows for a range of studies of planarian memory storage and inheritance. When trained planaria are sliced in half and allowed to regenerate, both the head- and tail-derived worms demonstrate significant recall of their original training. The tail-end fragments of planaria do not contain parts of the brain, implying that memory may be stored outside of the nervous system [3, 6].

In early studies of planaria, the worms were trained to run mazes, conditioned to form light/shock associations, and were even taught to crawl over a photoelectric beam to turn off a light [7]. In classical Pavlovian conditioning studies, planaria were taught an association between a flash of light and electric shock by systematically first flashing the light and then shocking the worms after a predetermined pause. It usually took several weeks for the planarians to learn; even then, results were weak because only a few worms could be trained at a time. Significant problems frequently arose in planarian memory studies, such as observer bias, small sample size due to the difficulty of training, and a lack of controls for sensitization and habituation effects. These weaknesses ultimately led to the abandonment of planaria studies [1, 7]. Despite this, researchers have managed to train planaria to a specified criterion of learning [4], which is defined as an alteration of an organism's original behavior as an adaptation to sensory stimuli in its surroundings."

Now ask yourself how even much more elaborate content discrimination and behavior might be possible with trillions of brain cells.

By Probeman (John Donovan)

Monroe wrote:
that doesn't help me...

I'll try more simply.

What is contentful about recognizing the pattern of a ripe red berry in a field of green leaves?

Monroe wrote:
It seems so far that this theory is incapable of accounting for (and inconsistent with) abstract thought, thinking about things not present or nonexistent, false representations, dreams... Any mental process that isn't directly about body-behavior.

And what about brain-stem functions? Does Dennett's theory have anything to say about why they are not conscious? Isn't there a distinct pattern of "winning" going on there with respect to digestion and heartbeat or whatever?

Absolutely. Those organisms that can't sort out homeostasis do not survive to reproduce. I suggest you read Damasio's "The Feeling of What Happens" which deals exactly with the integration of autonomic and limbic processes with both lower and higher levels of consciousness. He is a clinical psychologist whose ideas are very aligned with Dennett. Consider his work showing that those individuals that through brain injury lose connections to their limbic system also lose the ability to make rational decisions. We are intentional through limbic biochemicals.

There is an extremely large gap between simple homeostasis and abstract thinking but it is a continuum and that exactly why I want you to start with simpler behavior before going any further.

In fact this is normal reductionistic science. In order to understand something complex one first has to understand the simple things first. You want to jump from no brain to human brain in one step. This will ensure that you will not make any progress in your understanding until you work through it carefully instead of merely jumping around based on your intuitions.

By Probeman (John Donovan)

This is the abstract from the paper by Dennett that Faustus mentioned. It's worth reading and directly answers Descartes1's question.

"Theorists are converging from quite different quarters on a version of the global neuronal workspace model of consciousness, but there are residual confusions to be dissolved. In particular, theorists must resist the temptation to see global accessibility as the cause of consciousness (as if consciousness were some other, further condition); rather, it is consciousness. A useful metaphor for keeping this elusive idea in focus is that consciousness is rather like fame in the brain. It is not a privileged medium of representation, or an added property some states have; it is the very mutual accessibility that gives some informational states the powers that come with a subject's consciousness of that information. Like fame, consciousness is not a momentary condition, or a purely dispositional state, but rather a matter of actual influence over time. Theorists who take on the task of accounting for the aftermath that is critical for consciousness often appear to be leaving out the Subject of consciousness, when in fact they are providing an analysis of the Subject, a necessary component in any serious theory of consciousness. "

From :

<http://ase.tufts.edu/cogstud/papers/cognition.fin.htm>

By Probeman (John Donovan)

Descartes1 wrote: But that doesn't seem reasonable by the way they've gone about "[i]explaining[/i] consciousness."

Yes, if by "reasonable" you mean "intuitive" or "satisfying" (as Julios put it), then a scientific description of consciousness may fail you. Likewise many people today still can't intuitively accept that living tissue, like a kidney, is simply a collection of atoms. "Simply" is of course an understatement. The entire technological biochemical resources of the planet could probably not build a kidney atom by atom, much less a brain which is likely many orders of magnitude more complex.

The point I'm making however is that the desire to describe (or ascribe to) consciousness other metaphysical or religious properties or processes may be intuitively satisfying but that is not Dennett's goal. However, he wouldn't mind you to feel intuitively satisfied, hence his attempts to come up with metaphorically pleasing models like "multiple drafts" or "fame in the brain". But ultimately any explanation he offers has to answer to the scientific data. Once again there is no natural law that states that our scientific explanations have to be satisfying or intuitive. They can however become reasonable if one makes sufficient effort to understand them. It seems to me the onus is on (all of) you to do more reading rather than have Faustus and I spoon feed you the explanation.

Here is an extract from the paper that Faustus referred to that every one on this thread ought to read and re-read several times. Because I agree these ideas are unintuitive to be sure.

"A consensus may be emerging, but the seductiveness of the paths not taken is still potent, and part of my task here will be to diagnose some instances of backsliding and suggest therapeutic countermeasures. Of course those who still vehemently oppose this consensus will think it is I who needs therapy. These are difficult questions. Here is Dehaene and Naccache's short summary

of the global neuronal workspace model, to which I have attached some amplificatory notes on key terms, intended as friendly amendments to be elaborated in the rest of the paper:

At any given time, many modular (1) cerebral networks are active in parallel and process information in an unconscious manner. An information (2) becomes conscious, however, if the neural population that represents it is mobilized by top-down (3) 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 (4), 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 (5) what we subjectively experience as a conscious state. [from the ABSTRACT]

(1) Modularity comes in degrees and kinds; what is being stressed here is only that these are specialist networks with limited powers of information processing.

(2) There is no standard term for an event in the brain that carries information or content on some topic (e.g., information about color at a retinal location, information about a phoneme heard, information about the familiarity or novelty of other information currently being carried, etc.). Whenever some specialist network or smaller structure makes a discrimination, fixes some element of content, "an information" in their sense comes into existence. "Signal," "content-fixation," (Dennett, 1991), "micro-taking," (Dennett and Kinsbourne, 1992) "wordless narrative" (Damasio 1999), and "representation" (Jack and Shallice) are among the near-synonyms in use.

(3) We should be careful not to take the term "top-down" too literally. Since there is no single organizational summit to the brain, it means only that such attentional amplification is not just modulated "bottom-up" by features internal to the processing stream in which it rides, but also by sideways influences, from competitive, cooperative, collateral activities whose emergent net result is what we may lump together and call top-down influence. In an arena of opponent processes (as in a democracy) the "top" is distributed, not localized. Nevertheless, among the various competitive processes, there are important bifurcations or thresholds that can lead to strikingly different sequels, and it is these differences that best account for our pretheoretical intuitions about the difference between conscious and unconscious events in the mind. If we are careful, we can use "top-down" as an innocent allusion, exploiting a vivid fossil trace of a discarded Cartesian theory to mark the real differences that that theory misdescribed. (This will be elaborated in my discussion of Jack and Shallice below.)

(4) How long must this minimal duration be? Long enough to make the information available to a variety of processes-that's all. One should resist the temptation to imagine some other effect that needs to build up over time, because . . .

(5)The proposed consensual thesis is [b]not that this global availability causes some further effect or a different sort altogether-igniting the glow of conscious qualia, gaining entrance to the Cartesian Theater, or something like that-but that it is, all by itself, a conscious state.[/b] This is the hardest part of the thesis to understand and embrace. In fact, some who favor the rest of the consensus balk at this point and want to suppose that global availability must somehow kindle some special effect over and above the merely computational or functional competences such global availability ensures. Those who harbor this hunch are surrendering just when victory is at hand, I will argue, for these "merely functional" competences are the very competences that consciousness was supposed to enable.

Is this unreasonable? Still not satisfied? Then read the whole paper:

<http://ase.tufts.edu/cogstud/papers/cognition.fin.htm>
By Probeman (John Donovan)

[quote=Monroe]What if someone forgets a conscious state before they are able to report it, or chooses not to report it, or lacks the linguistic capacity to report it? The first thing happens all the time (i.e. dreams, when you know you were just thinking something but you can't remember what it was), the second happens a whole lot, and the third is theoretically possible and probably actual at some time or another and shows that they cannot be identical/equivalent conditions .[/quote]

Monroe, Dennett deals with these objections in detail in the Chap 4 post that Faustus and I wrote in the Dennett discussion thread. You can also read an update on this topic here:

<http://ase.tufts.edu/cogstud/papers/JCSarticle.pdf>

Here is a section from that paper that discusses your "unconscious consciousness" point:

[indent=20]"Consider masked priming. It has been demonstrated in hundreds of different experiments that if you present subjects with a 'priming' stimulus, such as a word or picture flashed briefly on a screen in front of the subject, followed very swiftly by a 'mask' — a blank or sometimes randomly patterned rectangle — before presenting the subjects with a 'target' stimulus to identify or otherwise respond to, there are conditions under which subjects will manifest behaviour that shows they have discriminated the priming stimulus, while they candidly and sincerely report that they were entirely unaware of any such stimulus. For instance, asked to complete the word stem fri___, subjects who have been shown the priming stimulus "cold" are more likely to comply with frigid and subjects who have been shown the priming stimulus "scared" are more likely to comply with fright or frightened, even though both groups of subjects claim not to have seen anything but first a blank rectangle followed by the target to be completed. Now are subjects to be trusted when they say that they were not conscious of the priming stimulus? There are apparently two ways theory can go here:

A. Subjects are conscious of the priming stimulus and then the mask makes them immediately forget this conscious experience, but it nevertheless influences their later performance on the target.

B. Subjects unconsciously extract information from the priming stimulus, which is prevented from 'reaching consciousness' by the mask.

Chalmers suggests that it is my 'attitude' that there is nothing to choose between these two hypotheses, but my point is different. It is open for scientific investigation to develop reasons for preferring one of these theoretical paths to the other, but at the outset, heterophenomenology is neutral, leaving the subject's hetero-phenomenological worlds bereft of any priming stimuli —that is how it seems to the subjects, after all — while postponing an answer to the question of how or why it seems thus to the subjects. Heterophenomenology is the beginning of a science of consciousness, not the end. It is the organization of the data, a catalogue of what must be explained, not itself an explanation or a theory. (This was the original meaning of 'phenomenology': a pretheoretical catalogue of the phenomena theory must account for.) And in maintaining this neutrality, it is actually doing justice to the first-person perspective, because you yourself, as a subject in a masked priming experiment, cannot discover anything in your experience that favours A or B. (If you think you can discover something —if you notice some glimmer of a hint in the experience, speak up! You're the subject, and you're supposed to tell it like it is. Don't mislead the experimenters by concealing something you discover in your experience. Maybe they've set the timing wrong for you. Let them know. But if they've done the experiment right, and you really find, so far as you can tell from your own first-person

perspective, that you were not conscious of any priming stimulus, then say so, and note that both A and B are still options between which you are powerless to offer any further evidence.)

But now suppose scientists look for a good reason to favour A or B and find it. What could it be? A theory that could provide a good reason would be one that is well-confirmed in other domains or contexts and that distinguishes, say, the sorts of discriminations that can be made unconsciously from the sorts that require consciousness. If in this case the masked discrimination was of a feature that in all other circumstances could only be discriminated by a conscious subject, this would be a (fairly) good reason for supposing that, however it may be with other discriminations, in this case the discrimination was conscious-and-then-forgotten, not unconscious.

Notice that if anything at all like this were discovered, and used as a ground for distinguishing A from B, it would be a triumph of third-person science, not due to anything that is accessible only to the subject's introspection. Subjects would learn for the first time that they were, or were not, conscious of these stimuli when they were taught the theory. It is the neutrality of heterophenomenology that permits such a question to be left open, pending further development of theory. And of course anyone proposing such a theory would have to have bootstrapped their way to their own proprietary understanding of what they meant by conscious and unconscious subjects, finding a consilience between our everyday assumptions about what we are conscious of and what we are not, on the one hand, and their own classificatory scheme on the other.

Anything too extreme ('It turns out on our theory that most people are conscious for only a few seconds a day, and nobody is conscious of sounds at all; hearing is entirely unconscious perception') will be rightly dismissed as an abuse of common understanding of the terms, but a theory that is predictively fecund and elegant can motivate substantial abandonment of this anchoring lore. Only when such a theory is in place will we be able, for the first time, to know what we mean when we talk about 'the experiences themselves' as distinct from what we each, subjectively, take our experiences to be. [/indent]

He also discusses it in the "Are we explaining consciousness yet?" article Faustus cited.

[quote=Monroe]If you're asking what the content of such a "recognition" is, then the content is of course, "hey, there's a red berry there in a field of green leaves." What is this supposed to prove?[/quote]

Ok, now ask yourself why this "recognition" might have been important during our evolutionary history when humans couldn't just walk down to their neighborhood supermarket and grab more food than they could possibly eat.

My long belabored point is that the "content" of such patterns is provided by the naturally selected "recognition" of our ancestors. Put simply, those that could distinguish the importance or usefulness of a pattern of a ripe red berry in an field of green leaves obtained a survival advantage over those that couldn't. By Probeman (John Donovan)

Hmm. Once Probeman posted the quote from the article, especially point five with the line in bold, I figured that Dennett's position would be clear, even if no one but Probe and I agreed with it. I'm almost at a loss, since I can't imagine what Probeman and I could do to make it more clear. (Though it must be noted that even Dennett says—in that very passage—that this is the hardest part of the theory for most people to grasp.)

Descartes I wrote:

If consciousness is not something other than winning processes, then why use terminology like "becomes conscious"? Why not just say "becomes dominant" and stop there?

Because the Multiple Draft theory and the workspace models are theories of consciousness, and it would be silly for them to leave the concept out!

The theories are there to explain what happens in your brain that enables you to remember and report informational content. In ordinary language, when you can remember something or report something, you are said to be conscious of it. So any scientific theory of consciousness will be one that explains the causal mechanisms which lead to the creations of such memories or capacity to report.

Perhaps the problem is (in part) due to a disconnect between ordinary language (folk psychology) and scientific language. They are there to do different work. The folk psychology of consciousness leads us to think of it in terms of a special property that something acquires. The scientific models discussed in this thread do not, but rather define conscious states in terms of work being done, capacities enabled, and so forth. So when the scientific model says, "And once a content gains global accessibility (or dominance), it becomes conscious," it sounds as if the model is describing a transformation of some sort, or the addition of some extra property. But rather, the model is simply saying that once this event occurs, that various capacities are enabled in the subject. These are the capacities which, in ordinary language, lead to situations where we might say that someone is conscious of X, where X is the content that just "won" the internal Darwinian competition for access.

By Faustus (Brian Peterson)

other-wise wrote:

- 1) **Dennett** defines consciousness as: the state of remembering and/or reporting information.
- 2) When information processed by a given area(s) of the brain becomes globally available to other brain processes, said state occurs.
- 3) Information that is dominant (AKA "globally available") does not *become* conscious; dominance *is* consciousness.

Did I miss (or mis-characterize) anything?

Faustus should chime in himself, but the only thing I would say is that point number one is clearly incomplete. Remembering and reporting are certainly a large part of conscious behavior but not all of it. Dennett offers this poetic description of consciousness:

"There is no single, definitive "stream of consciousness," because there is no central Headquarters, no Cartesian Theater where "it all comes together" for the perusal of a Central Meaner. Instead of such a single stream (however wide), there are multiple channels in which specialist circuits try, in parallel pandemoniums to do their various things, creating Multiple Drafts as they go. Most of these fragmentary drafts of "narrative" play short-lived roles in the modulation of current activity but some get promoted to further functional roles, in swift succession, by the activity of a virtual machine in the brain. The seriality of this machine (its "von Neumannesque" character) is not a "hard-wired" design feature, but rather the upshot of a succession of coalitions of these specialists.

The basic specialists are part of our animal heritage. They were not developed to perform peculiarly human actions, such as reading and writing, but ducking, predator-avoiding, face-recognizing, grasping throwing, berry-picking, and other essential tasks. They are often opportunistically enlisted in new roles, for which their native talents more or less suit them. The result is not bedlam only because the trends that are imposed on all this activity are themselves the product of [evolutionary] design. Some of this design is innate, and is shared with other animals. But it is augmented, and sometimes even overwhelmed in importance, by microhabits of thought that are developed in the individual, partly idiosyncratic results of self-exploration and partly the predesigned gifts of culture. Thousands of memes, mostly borne by language, but also by wordless "images" and other data structures, take up residence in an individual brain, shaping its tendencies

and thereby turning it into a mind.”

I would also add that point number two should mention the internal competition for dominance by various multiple drafts. By the way, this concept helps to explain "slips of the tongue" and other mis-utterances.
By Probeman (John Donovan)

Yo, Other-wise! Where ya been all this time?

I'll go over your three points one at a time.

Other-wise wrote:

1)Dennett defines consciousness as: the state of remembering and/or reporting information.

There's more to it than that, per Probeman's post, but I suppose that would do as a sort of cartoon picture. Maybe I'd amend the cartoon to something like "a dispositional state where information content can be remembered, reported, or otherwise acted upon." The point is that a functionalist like Dennett is always going to define consciousness in terms of stuff it enables you to do, where stuff it enables you to do pretty much exhausts the concept (i.e., there aren't qualities or properties of consciousness that are enabling, and then FURTHER qualities or properties that count as the qualia).

Other-wise wrote:

2) Information processed by a given area(s) of the brain competes for global availability to other brain processes. When an information process becomes globally available, said state occurs.

Can't think of any amendments at this time, so I'll go with this!

Other-wise wrote:

3)Information that is dominant (AKA "globally available") does not become conscious; dominance is consciousness.

I'm sort of mixed on this one. Part of me wants to say, "Okay as a cartoon or sound-bite," but I'm less comfortable with this summation than I am with point one.

Part of the reason is something I alluded to in my recent reply to Descartes1, about the relationship between vocabularies of folk psychology and those of science. It doesn't come out in the interview or (I think) any of the cited articles, but Dennett doesn't believe in a reductionist mapping of one to the other. So what I'm wary about is the interpretation one might give to "is". So long as the identity is thought of as fuzzy and casual rather than strict (as in "heat is molecular motion"), I'm fine with your way of putting it.
By Faustus (Brian Peterson)

other-wise wrote:

("Remembering" or "access" should probably be in there someplace, but since they tend to imply that *something* is accessing, or *something* is remembering...

Yes, something is- the brain!

By Probeman (John Donovan)

other-wise wrote:

probeman, my first response to this was:

Just trying to avoid any whiff of homuncularity, which we both know Dennett would abhor (Note that I specifically invoked the brain... twice... in the very first statement... as the locus of action).

Other-wise,

My first response to this is: exactly! There is no "you" inside "you." And remember Dennett would say that homunculi are perfectly acceptable as explanation provided that the homunculi used to explain other homunculi must be less conscious/intelligent than the homunculi they are explaining.

other-wise wrote:

Is "remembering" or "accessing" done by the *brain* or a brain *process*? That may seem like a nitpick, but it's exactly the type of thing that can leave people confused and argumentative.

One could say it was nit-picking except that I agree this semantic shift between the folk and scientific usages seems to be an endless source of confusion for both lay people and even scientists who should know better. Though it might be less of a problem if there weren't still those that insist upon essentialistic metaphysically intuitive explanations. That's why I suggest to people that they read Dennett. Because he deals so forthrightly with the unintuitive nature of the problem from a philosophical perspective.

I would say that the "sub work" is done by "sub processes", but the work as a whole (consciousness) is done by the whole process (brain). It's clear enough from the last 10 years of cognitive science that there are many specialized sub-processes in the brain that contribute to and/or manifest the behavioral aspects that we externally recognize as a conscious mind. Even before brain imaging techniques these sub processes were visible from studies of brain injured patients. What is being worked out now is how these brain processes interact globally. The idea of a competitive yet cooperative interaction between these processes inside our skulls offers some possibilities for explaining what we actually observe.

And this also might explain why first person observations are unreliable for this purpose- since those "observations" are simply those that "became famous" and therefore are not in a position to describe that process since the process itself was almost all preconscious. Did you read the Dennett paper Faustus suggested "Are We Explaining Consciousness Yet?"

By Probeman (John Donovan)

Descartes1 wrote:

I'm not confused as to *what* Dennett's position is. I'm wondering about how it makes sense.

Maybe you're not confused about "*what* Dennett's position is" but you clearly haven't read enough to understand the implications of his position. Your need for it to make "sense" is fine if you are asking how a scientific explanation is possible. But from your following paragraph I rather suspect that instead you are demanding that Dennett's explanation to be intuitively obvious. You say you are "dropping all intuitions", but then you immediately proceed to invoke them. Why should any theory of science be required to conform to our "subjective experience"? It is my "subjective experience" that the Sun circles the Earth and that living organisms must contain some vital substance more than the mere arrangement of their atoms. So? My subjective experiences in these regards are almost certainly wrong.

It is no different for the scientific study of the human mind. I know it seems to all of us that "there surely is something about consciousness of the visual data itself that is different from neural activity". But this is merely our often fallible intuitions speaking, not science. We all have the built in naive explanation that

there is something more- yet there is no evidence for this claim. What you should know already from cognitive science is that "how it seems to us" is simply not "how it really is". The all too human mistake is to assume that because it is our mind that we are introspecting, it must therefore be an incorrigible source of explanation or understanding.

I really think that if you are serious about having this "make sense" then you need to make a concerted effort to read more detailed works on the subject. Have you at least read the "Are we explaining consciousness yet?" paper? That would be a good place to start. Also try reading the Dennett Discussion chapter summary threads that Faustus and I wrote. I will bring them forward for you.

By Probeman (John Donovan)

Descartes I wrote:

Instead of that, however, Dennett just denies that it even happens. It doesn't seem to me that I'm invoking any intuitions, but I cannot deny what is evidently clear to me; that is, I cannot deny that, say, the actual feeling of the pain has a qualitative property to it when I experience it. And this qualitative property is evident every time I, say, pinch myself. To deny this is impossible simply because I am constantly reminded of the contrary every time I feel pain.

Dennett explicitly does not deny that people have experiences. What he denies is that experiences can only be understood as having all the qualities and properties that pre-scientific philosophical tradition assigns to them. The chapter four summary on heterophenomenology is where he begins to spell out what this could mean. Unfortunately, the book doesn't really take off until he starts his attack on qualia, and Probeman and I haven't made it that far in our summary. (You can blame me—I'm months late delivering my chapter ten summary.)

Descartes I wrote:

But I cannot fathom how that very scientist could also know the pain and experience it in the same qualitative way that I do, while I am experiencing it and he is just studying my brain's neural activity.

No, neither can I. I can't even fathom why you would think Dennett is claiming such a thing! I could be mistaking you on this point, but it's almost as if you are expecting complete theory of how the brain produces sensation to produce in its practitioners the very sensations they are studying. To me that's like suggesting a theory of hurricanes would be wet and windy. I must be misunderstanding to you a very extreme degree, but at this point I can't really get what you actually meant.

One way I would put things is that when we talk of a subject's experiences in a scientific context, we are talking about informational states of brain tissue being represented and reported upon by other informational states of brain tissue in the same brain. (See the discussion of Shakey in the chapter four summary when he tries to tell us how he tells pyramid shapes from boxes.) That's one way of accessing the content of the experience. Another is to do a brain scan and extrapolate from that using whatever knowledge of neuroscience we've managed to put together. Since the scientist is getting the information in a completely different way, she would not be expected to have an identical experience.

[xeno wrote:

Julios]I too need to read his stuff, but it seems to me that he is metaphysically sandwiching "experience" into "substance/process", and then claiming it to be explained.

Is "experience" perhaps not a candidate for a metaphysical category of being?

Is me asking this very question missing Dennett's point?

Not sure what you mean, here. Personally, I tend to favor the view that introducing metaphysics into

consciousness studies is the source of a lot of problems, so I would be inclined to NOT understand experiences on the model of a metaphysical category.
By Faustus (Brian Peterson)