

COMMENT ON VICENTE AND JORBA

Daniel Dennett

Marta Jorba and Agustín Vicente (hereafter, V&J) raise questions about how my view meshes or clashes with the position of Thomas Metzinger. Yes, Metzinger and I are on very similar paths. I owe a discussion of our agreements and residual differences, but that will have to await another occasion.

V&J doubt my claims about the illusory nature of human consciousness, and provide a detailed set of objections to what they take me to be claiming. And it is clear from what they say that I have not done as good a job as I should have on clarifying my own positions. They often note, in fact, that they are unsure whether they have interpreted me correctly (since some of my views seem, to them, well-nigh incredible), and sure enough, *some* of the views they have their doubts about are not what I meant. Let me go through a few of their cases to try to set the record straight, while actually using their sympathetic probing to refine my own thinking.

First let's look at reasons. In their footnote 4 they note:

At some points, he seems to hold that we can actually know reasons. At some other points, he seems to argue that everything we think we know, including reasons, is illusory knowledge, which can be more or less approximate.

Yes, I do think we can actually know our reasons, and maybe we even *usually* know them, but on these occasions our authority is only circumstantially better than that of our companions and interlocutors. We know our reasons because we can recall (at least for a brief period) our thinking. The point is delicate. Imagine encountering a friend who is busily measuring and sawing boards, and asking "what are you doing?" And imagine your friend replying: "What a good question! You know, I should ask my analyst the next time I have an appointment. I'm sure

there is some deep purpose to my current behavior. If only I could discover it!" You would probably conclude that your friend was playing a joke on you. We all know (and know that we all know) that in general, when we set about doing some fairly complicated thing, or even some routine and clearly purposive thing, we have a ready and reliable answer to the question "what are you doing?" We wouldn't ask, otherwise. People are just able — as Anscombe (1957) noted — to give their reasons, right off, and it is only in rather unusual circumstances that they draw a blank. As we grow older, of course, we are apt to find ourselves walking into the kitchen and wondering why we are doing this, having forgotten the mundane errand we had set out on, but these are embarrassing cases of forgetfulness, not — except in extreme cases — utter absence of self-knowledge, as we typically prove to ourselves by recovering the temporarily lost quest after a few seconds of pondering. But notice that if we were in the habit of talking aloud to ourselves as we pattered around, our companions could have almost as good evidence as we do about what we were doing and why. Standing perplexed in front of the coat closet, I look to my wife who says "You just said you were going to get your windbreaker."

When we know what we've said to ourselves, silently or aloud, we have reliable but not infallible or incorrigible knowledge of our reasons. When we don't talk to ourselves or even reflect wordlessly on what we are doing, we "know what we are doing and why" only the way bears and birds know what they are doing and why: there are control systems engaged that are (well-)designed to modulate action by feedback and anticipation so that proximal behavioral goals are achieved, barring mishap. The reasons then may have been articulated at some point (by us) and then set aside, or they may be like the reasons of animals, free-floating rationales never expressed or thought about until curious and imaginative human interpreters came along.

When V&J turn to "phenomenal" experience, they raise some good questions about my position:

There seems to be a certain conceptual tension in maintaining that the field of appearance itself — experience — is an illusion: in consciousness (in this sense), what appears to be is what it is.

Let's see if I can dissolve this conceptual tension. I do not know just what they are packing into the idea of "the field of experience" but let me try to illustrate it with an example. When I open my eyes in daylight,

I can readily report, search for, track, describe all manner of things, and if asked how I am able to do this, I reply “I see them, of course.” But if asked what that involves, what happens in between the photons falling on the retinas and the confident, detailed, visually informed activities I then engage in, I may theorize, but introspection tells me nothing more. I have access to contents (“I can’t make out the writing on the distant sign, but it’s in red capital letters”) but not processes and I know nothing “from the inside” of the medium.

Richard Power (2017) has put the point very well:

We understand the concept of representation from external representations, such as pictures, or verbal descriptions. For these representations we can have direct experience of both a representer (e.g., portrait painting) and a representee (e.g., the person painted). Call these the medium and the content. . . . In short, we conceptualise the medium of our internal representations by abstracting some features from the content, and attributing them to some kind of spiritual or ghostly substance. That is the best we can do, since actually we cannot experience the medium at all and have to look for analogies in the external world.

This is the conceptual scheme that we bring to internal representations, because it is the only one we have. But there is a huge difference. For external representations we can experience both medium and content, oil on canvas as well as people, trees, or whatever. But for internal representations, we do not experience the medium AT ALL. The idea that the medium is some state of the brain seems intuitively absurd, so powerful is the illusion that we are dealing with an iconic representation in a medium of spirit.

So the illusion Power is pointing out is an intellectual illusion that is generated by human beings reflecting on the task of explaining to themselves how they can have experience. It is extremely unlikely that any other species is capable of such curiosity, so they are immune to *this* illusion, but, like us, they are the *beneficiaries* (not the victims) of species-specific *user-illusions* that simplify the world they perceive and act in. Chimpanzees and dolphins are neither dualists nor materialists, but, like us and all other perceiving animals, they *employ* versions of the physical world that “create” properties, such as colors, and aromas, that exist only relative to the property-detection equipment of the organisms. These are “functional” properties of things in the world, not properties of renderings in the head, and the manifold of reactions by the organism to these “represented” properties exhaust the what-it-is-like-ness or subjectivity of that organism. The “intrinsic” properties of the representations (as

contrasted with the properties represented by those representations) are whatever “intrinsic” properties neural signals, spike trains, reverberations have; it is the functional properties of these physical mechanisms and the physical features of the world they track that matter for creating the idiosyncratic subjectivity of the organism. The sweetness of honey is not an intrinsic property of glucose or an intrinsic property of some internal rendering of honey, any more than the funniness of a joke is an intrinsic property of the string of words expressing it or of the excitation of one’s “funny bone” or “sense of humor” in the brain [Hurley, Dennett and Adams, (2011)].

Notice that the verb I italicized in the previous paragraph was “employ,” not “enjoy.” V&J say in footnote 5:

... it seems that for Dennett having an experience requires having a previous judgment that one is having that experience: “It is like something to be you because you have been able to tell what’s like to be you!” [p. 344]. So, if language is required in order to have that judgement, then language would be necessary to have experiences.

This is almost right, but I am not saying that one has to have a “previous judgment” (with or without the help of language) in order to have experience in the sense of being awake and informed (or misinformed) by one’s sensory apparatus. I am saying that since non-human animals don’t have the (language-borne) reflective capacities we have, there is an important sense in which a bat doesn’t know what it is like to be a bat! Their immunity to the Power illusion is also their ignorance of what it is like to be them! We can say, if we feel the urge, that of course bats know what it’s like to be bats, and fish know what it’s like to be fish, and so forth, but the only sense that can be given to these assertions is that these creatures are well-served by their cognitive equipment, which they exploit by “second nature” (not having to learn how to use it). They can’t compare notes, discuss it with their young, write novels about it, so their knowing is only *knowhow*. Can one have experience without *knowing that* one is having experience? I doubt that there is a good case to be made either way, since our everyday concepts of experience and knowledge are so porous and labile.

A good reason to *avoid* avowing that (“surely”) bats know what it’s like to be bats is that it tends to create an unanswerable question, a quixotic demand for science to *draw the line* between species that *enjoy* “phenomenal consciousness” and those that don’t. The problem with using

the verb “enjoy” is that it perpetuates the postponement of the Hard Question: “And then what happens?” [Dennett (1991)]. It puts “the Subject” (the enjoyer, the appreciator, the sufferer) into the *next* chapter of the book explaining consciousness *and then never attempts to write that chapter!* [Dennett (forthcoming)]. This postponement is perfectly illustrated by a passage, with which I otherwise fully agree, about pain experience:

This pain experience, presumably a user-friendly representation, tells *the user* that the fibula is wrong, that she has to do something with her fibula, or both things at the same time, for instance. In this, the experience plays the same role that the underlying reality plays, for neurons can also be said to (ultimately) represent the damage and to trigger the corresponding behaviour. If this is right, then it seems that a certain experience can tell *us* something about what certain physical tokens in the brain represent. That is, the experience does not tell *us* what really goes on in our brains (the neural processes) but it can tell *us* that whatever goes on in our brains, it represents a certain content. [italics added—DCD]

Who is this *us*? There is no isolated single *user* of all the illusions in the user-friendly system. The consumers, as Ruth Millikan would put it, of these representations are themselves subsystems of the whole cognitive apparatus. If you think that in addition to all those subsystems there is one special subsystem, the *I* that is being “told” all this (in what medium — “phenomenal” pain language?) then you are avoiding the Hard Question. [See also Huebner and Dennett (2009)]. V&J say:

We want to argue that we can know the content of our inner speech episodes also by introspection. That is: an educated observation of what goes on in our minds can tell *us* what our brains are representing. That is, in principle, *we* can get full introspective access to the content of our brain representations. It can even be argued that being wrong about the content of an inner speech episode is more difficult than being wrong about the features of the vehicle. In overt speech at least one usually knows what she herself meant, although she may be wrong about what she actually said – or how she actually said what she said.

Yes, we can improve with practice the activity of “introspection” but to say that “*we* can get full introspective access to the content of our brain representations” is to divorce *us* from our brains. In one sense, almost trivially, our brain systems have access to the content of our brains representations — that is what makes them representations with the content they have. (In a similar vein, computer programs have full access *by definition*

to all the content of their representations; they may have “myopic” access, but then the content of their representations is myopic. A chess playing program may not know it is playing chess, but it has full access to the content of its representations. In the other sense in which our brains have content, for *me* to gain access to (some of) the content is quite an achievement; when small children can start telling us what they know about the content of their minds, they become experiencers who know they are experiencers, and know what it’s like to be them.

Finally, consider the last sentence in the quotation, which has a slippery “know” in it. Does she *know* what she meant if she can’t *tell us* what she meant? Students often claim to know what they meant by a sentence in an essay and then discover, to their frustration and embarrassment, that they don’t really know what they meant; it just seemed to them as they wrote it that they knew what they meant. The authors do say “usually” and that saves the sentence as true, I think, but not if it is read as asserting a normal introspective authority that is more than circumstantially better than that of the second-person interlocutor. See Dennett (1996), “The Case for Rorts” for an extended discussion of this situation. I am grateful to V&J for raising these issues, since it gives me an opportunity to resurrect that largely ignored essay, which has, I think, a number of points that bear well on current issues.

REFERENCES

- ANSCOMBE, G. E. M. (1957), *Intention*, Oxford, Blackwells.
- DENNETT, D. C. (1991), *Consciousness Explained*, New York: Little Brown.
- 2000, “The Case for Rorts,” in Robert Brandom, ed., *Rorty and His Critics*, Wiley-Blackwells.
- 2018, “Facing up to the Hard Question of Consciousness”; in *Royal Soc. Philosophical Transactions B*, guest edited by Peter Fazekas and Morten Overgaard. 30 July 2018. DOI: 10.1098/rstb.2017.0342
- HUEBNER, B., and D. DENNETT, (2009), “Banishing “I” and “We” from Accounts of Metacognition,” with Bryce Huebner, *Behavioral and Brain Sciences*, 32:2, pp. 148-9.
- HURLEY, M., DENNETT, D. C. and Adams, R. B., (2011), *Inside Jokes: Using Humor to Reverse-Engineer the Mind*, MIT Press.
- POWER, R., (2017), personal correspondence.