

Journal of Cognition and Neuroethics

Libet, Free Will, and Conscious Awareness

Janet Levin

University of Southern California

Biography

Janet Levin is Associate Professor of Philosophy at USC. Her research interests include philosophy of mind, philosophy of psychology, and epistemology.

Publication Details

Journal of Cognition and Neuroethics (ISSN: 2166-5087). March, 2015. Volume 3, Issue 1.

Citation

Levin, Janet. 2015. "Libet, Free Will, and Conscious Awareness." *Journal of Cognition and Neuroethics* 3 (1): 265–280.

Libet, Free Will, and Conscious Awareness

Janet Levin

Abstract

In a series of well-known experiments, Benjamin Libet fits subjects with electrodes that monitor their brain activity, and instructs them to decide whether or not to flex their wrists at various times during a certain interval and then follow through. And—notoriously—he finds that the subjects' wrist flexings are preceded by the occurrence of a 'readiness potential' (RP) that begins about 400ms *before they report any conscious inclination or wish to act*. Therefore, Libet argues, these wrist flexings do not arise from the subjects' free will. There have been numerous attempts to dispute Libet, and argue that his subjects' conscious wishes or inclinations can be regarded as the causes of their actions—and I find many of these arguments compelling. Here, however, I question the connection between conscious motivation and freedom of action, and argue that behavior produced by wishes or inclinations of which we are not consciously aware can often be viewed as sufficiently up to us, or under our control, to count as free action. On the other hand, as I argue, we may need to be consciously aware of our motivations to be held *morally responsible* for what we do. And this, I suggest, has some potentially interesting implications for our common views about the relation between free will and moral responsibility.

Keywords

Libet, free will, conscious experience, moral responsibility

In his well-known experiments that purport to show that we have less free will than we may think, Benjamin Libet (1985, 2011) fits subjects with electrodes that monitor their brain activity, and instructs them to decide whether or not to flex their wrists at various times during a certain interval and then follow through. And—notoriously—he finds that the subjects' wrist flexings are preceded by the occurrence of a 'readiness potential' (RP) that begins about 400ms *before they report any conscious inclination or wish to act*. Therefore, Libet argues, these wrist flexings do not arise from the subjects' free will. It may *seem* to the subjects that they are consciously willing to flex their wrists, but this is merely an illusion.

On the other hand, Libet argues, the data show that if the subjects become aware of their inclinations to flex and consciously 'veto' such inclinations, their subsequent actions *are* directly initiated by a conscious process, and thereby *do* arise from their free will—at least insofar as their conscious vetoes are not themselves determined.¹ In short, on Libet's

1. As V. S. Ramachandran (and subsequently many others) have put it, for Libet there is no freedom of will, but only freedom of 'won't.' See, however, Lau et al (2007) for skepticism about whether vetoing is a

view, being initiated by a conscious wish or willing is *necessary* for an action to originate from a subject's free will—though perhaps not *sufficient*.

There have been numerous attempts to show that, contrary to Libet's suggestions, a subject's conscious wish or inclination to flex *can* be regarded as the cause of the flexing. For example, some (e.g., Roskies, 2011, Dennett, 2003, and Mele, 2011) argue that the activation of the RP is merely the *lead-up* to a subject's conscious willing, and not the willing itself; others (e.g., Roskies, 2011) suggest that the activation of the RP may well *be* the subject's conscious willing, which precedes (by a few ms) the subject's *report* (or even *conscious awareness*) that it has occurred. Alternatively, Horgan (2011) argues that, even if the initiation of the RP truly precedes the conscious willing to act (and not just the conscious awareness of that wish or inclination), that conscious state can be regarded as the *sustaining cause* of the implementation of a standing intention to act, and thus does not threaten the veridicality of the experience of conscious will.² I find many of these arguments compelling. However, they all focus on challenging the claim that the actions in question do not originate from the agent's *conscious wish or inclination to act*—and this implies that being initiated by a conscious willing is *necessary* for an action to be free.

Here, however, I want to question the connection between conscious motivation and freedom of action, and consider whether behavior produced by wishes or inclinations of which we are *not* consciously aware can nonetheless be viewed as sufficiently up to us, or under our control, to count as free action. I will argue that the answer, at least sometimes, may be 'yes,' and thus that the focus on whether Libet's findings undermine the view that our actions are produced by conscious motivation may be less relevant to determining whether we have free will than is often assumed.

On the other hand, I also will consider the relation between conscious motivation and *moral responsibility*; in particular, whether (we think) an agent can be blamed or praised for doing something if she has no conscious awareness of a decision, or wish, or inclination to do so. And I will argue that in this case the answer, more often, may be 'no.' The upshot of these considerations will be that there is reason to think that although an agent's conscious decision (or wish or inclination) to do A may not be necessary for

conscious activity, and Mele (2013c) for discussion of these, and other, findings.

2. In a different attempt to counter Libet, still others (e.g., Roskies, 2011, Mele 2013a,b, 2014) argue that even if Libet is correct to claim that his subjects' wrist-flexings do not arise from free will, this conclusion may not generalize to actions that are the products of more extensive deliberations in which conscious decision plays an important *contributing* role.

A to be *free*, (we tend to think) it may be necessary for the agent to be held *morally responsible* for doing A. And this has some potentially interesting implications for our common views about the relation between freedom of action and moral responsibility.

To think about these issues, let's consider a more 'real life' scenario that, at least at first glance, has affinities with Libet's wrist-flexing experiments:

Suppose I'm a professional diver, and I practice for at least three hours every day. I go up to the top of the high dive—without thinking much about what I'm doing—position myself, and dive directly into the water. I swim to the edge of the pool, then climb up and do it again—and again. Occasionally, however, something doesn't seem quite right: some debris in the water, a kid who's swimming toward the diving board, or something I can't quite put my finger on—and I don't make the dive, or I dive in a different direction. When everything is going well, it doesn't seem like I'm *consciously willing* to dive (or to dive in the particular direction that I do) or even that I'm consciously aware of an inclination to do so. Maybe, given that I told my coach that I'd indicate when I would attempt a dive, I give a thumbs up as I'm about to leave the platform. But sometimes I forget and just do it.

This vignette, of course, has at least some commonalities with the situation of Libet's subjects: when things go well, my indication of intention (thumbs up) seems equally after the fact. And I suspect that if there had been (waterproof) electrodes affixed to my head, my brain activity would look similar to that of Libet's subjects.³

As we know, Libet contends that in the situations in which his subjects decide to flex their wrists and do so, they don't act freely; only acts initiated by a conscious veto can be the products of free will. But is it so obvious that when I make a straightforward dive in the situation described I don't act of my own free will? I suspect that this may seem less clear. And this is so, it seems, even though my diving—as in many other cases of so-called *skill exercise*—may seem *more* 'automatic,' *less* governed by anything like conscious will, than the wrist flexing of Libet's subjects. Indeed, it's not clear that my diving in a different direction (or not diving at all) when things seem sketchy is best regarded as the result of a 'conscious veto' of a wish or inclination; these actions seem pretty automatic as well, at least if I'm truly a skilled diver whose training has made

3. After all, when I do take the dive, there must be *something* going on in my brain that precedes my action prior to my giving a thumbs-up to my coach, just as there is something going on in the brains of Libet's subjects before they express their inclination to flex their wrists.

for flexibility of response. But here too—or so it seems—my action (or refraining from action) seems to be something that I did freely, something that was up to me.⁴

Libet, no doubt, would disagree. But is it clear that this is the right verdict? In what follows, I want to address three questions. First, are there conditions under which seemingly automatic actions like diving—actions not obviously caused by a conscious decision—can nonetheless be products of free will; second, do these conditions require that the agent have *any* sort of conscious awareness—and if so, awareness of *what*; and third, do we have the same views about the relation between conscious awareness and moral responsibility?

First, let's examine the relation between freedom and automaticity. Some theorists argue that automaticity—even the sort exemplified in skill exercise like piano playing or diving—is incompatible not only with *free* agency, but with any sort of intentional agency at all.⁵ But this view is far from universal. Among the dissenters is Mele (2011), who argues that intentional action can occur without a conscious decision to act in that way when the action is routine, and suggests (26) that if subjects have a *conditional intention* to do something when they feel like it, and act forthwith, then their actions can count as intentional. Wayne Wu (2013a) goes further, and argues (258) that, in 'normal action' produced by intention, 'intentions are persisting nonphenomenal states of subjects that coordinate and constrain one's meandering through behavioral space.' I find these views about agency plausible, and if they are right, then it seems that my diving and Libet's subjects' flexings can at least sometimes be intentional actions even if not produced by an explicit conscious decision. But I want to go even further and question whether an act can be intentional not just if the agent has no conscious intention, conditional or not, to perform that act, but whether conscious awareness of *anything* at all is required for agency.

-
4. There are, of course, many examples of (what seems to be) free action without conscious choosing that aren't examples of skill exercise, such as pulling out one's wallet to pay the check, or getting up to answer the ringing doorbell. See Vihvelin (2013, 6.3) for other good examples of choice without conscious choosing, and discussion.
 5. See, for example, G. Strawson's (2003) argument that automaticity undermines agency. Also, consider Nadelhoffer (2011, 178), 'The agential threats that we will be examining here are ultimately fueled by the fact that the conscious mind exercises less control over our behavior than we have traditionally assumed. It is this deflationary view of conscious volition that is potentially agency undermining.' He continues (183), 'So whereas Libet's view [involving the possibility of 'veto power'] merely shrinks the domains over which we exercise control, Wegner seemingly leaves the conscious mind out of the causal loop altogether...As such, Wegner's view...represents a global and not merely partial agential threat.'

Now one may think that this suggestion is crazy: for an individual to perform a bona-fide intentional action, one may think, she must at least have conscious awareness of *something*, perhaps the environmental conditions in which she does what she does. After all, how could a piano player or a diver possibly play the right notes or dive in the right direction if she were not consciously aware of such things as the location of the piano keys or the water?

But is this connection so clear? Consider, for example, David Chalmers's (1995) distinction between the 'hard' and the 'easy' problems of consciousness. The hard problem, according to Chalmers, is to give a satisfying explanation of why it is that 'when our cognitive systems engage in visual and auditory information-processing, we have visual or auditory experience: the quality of deep blue, the sensation of middle C...[that is] why there is something it is like to entertain a mental image, or to experience an emotion.' This problem, he contends, is different from the 'easy' problems of providing a satisfying explanation of phenomena such as:

- (1) the ability to discriminate, categorize, and react to environmental stimuli
- (2) the integration of information by a cognitive system
- (3) the reportability of mental states
- (4) the ability of a system to access its own internal states
- (5) the deliberate control of behavior
- (6) the difference between wakefulness and sleep.

These problems, Chalmers acknowledges, are hard, but hard in a different way, in that even though it may take time and effort to solve them, there is 'no real issue about whether these phenomena can be explained scientifically. All of them are straightforwardly vulnerable to explanation in terms of computational or neural mechanisms.'⁶ And the reason is that it seems that all we have to do to provide a satisfying explanation of those phenomena is to (i) get clear about the way these processes function—and then (ii) look around (in the brain and body) for some mechanism that performs or implements this function. Once we find such mechanisms, the question seems closed; we're done.

6. He continues, 'To explain access and reportability, for example, we need only specify the mechanism by which information about internal states is retrieved and made available for verbal report. To explain the integration of information, we need only exhibit mechanisms by which information is brought together and exploited by later processes. For an account of sleep and wakefulness, an appropriate neurophysiological account of the processes responsible for organisms' contrasting behavior in those states will suffice. In each case, an appropriate cognitive or neurophysiological model can clearly do the explanatory work.'

But this is not so, Chalmers contends, for the how and why of conscious experience. Even if we were to have a comprehensive scientific explanation of (1)-(6), we would have no satisfying explanation of why a creature with these capacities has the conscious experiences it has—or even why it has any conscious experiences at all. Indeed, he argues, it seems possible that there could be creatures just like ourselves with respect to (1)-(6) but with no phenomenally conscious experiences.⁷

Now if the idea of such a creature is coherent, then it's conceivable that an individual could perform a wide variety of behaviors produced by sophisticated cognitive processes even though it doesn't have conscious awareness of environmental conditions—let alone any *experience of agency* that some take to a necessary antecedent of an intentional action. And this suggests that it's at least conceivable that a creature could be an intentional agent without being conscious. Indeed, we don't have to invoke such Chalmersian constructs to make the point, since there is at least some empirical evidence (from Kentridge et al, 1999) that the well-known blindsight patient GY can discriminate objects in his blind hemi-field. If this is so, then conscious awareness of one's environment may be merely a *de facto*, but not *in principle*, requirement for intentional agency.⁸

-
7. For example, he continues, 'we can imagine that right now I am gazing out the window, experiencing some nice green sensations from seeing the trees outside, having pleasant taste sensations through munching on a chocolate bar, and feeling a dull aching sensation in my right shoulder' and has a counterpart wrt (1)-(6) that 'will be perceiving the trees outside in the functional sense, and tasting the chocolate in the psychological sense, [and will be] awake, able to report the contents of his internal states, able to focus attention on various places...[but] none of this functioning will be accompanied by any real conscious experience. There will be no phenomenal feel' (1996, 94-5). Now, I have characterized Chalmers as arguing that it is imaginable that there are creatures that possesses capacities (1)-(6), but have no conscious experiences. But he makes a considerably stronger claim, namely that we can conceive or imagine creatures that are *our exact molecular duplicates* but have no conscious experiences. This claim is more controversial than the one I am considering, but for the purposes of this discussion the weaker claim is strong enough. (In their discussion of 'the zombie challenge' in the introduction to their (2013) collection, Vierkant, Kiverstein, and Clark are clear about this distinction.)
8. See Wu (2013a), who discusses these findings in service of his view that attention can be 'selection for action' even if not conscious: 'we would fail to fully capture an essential psychological capacity were we to restrict talk of attention to just these conscious forms. For the very capacity for action requires that the agent exhibit a striking form of attunement to the world so as to guide her behavior, and much of this attunement is in a way subterranean to consciousness even if it is not subpersonal. Responsiveness to the world, in action, precisely involves a way of attending to the world, more often unconscious than not.' See also Wu (2013b).

Nonetheless, an action can be intentional without being free, and so we need to consider whether *free* intentional action may require some sort of conscious awareness.⁹ In the literature on free will and free agency, theorists have suggested a variety of conditions that must be met for an action to arise from true freedom of the will. In much of this literature the guiding question is whether these conditions could be met in a deterministic universe. But I want to ignore that question and focus solely on whether one can meet any of these conditions without conscious awareness.

Consider, first, Harry Frankfurt's well-known (1971) account of what it takes for an act to proceed from a 'will that is free,' namely, that (i) I act freely; that is, *I do what I want*, and (ii) I act of my own free will; that is, I *want* to act (or I approve of my acting) on that desire, and finally, (iii) *if I hadn't wanted to act on that desire, it wouldn't have been the one that caused my action*.

This characterization, it seems, reflects some intuitive distinctions we may be inclined to make in variations of the diving case. Suppose I'm preparing to dive, as usual, and a little kid swims under the board—and I make my usual dive and hit him. Did I do *this* of my own free will? Well, it seems that we would have to consider a number of things: Did I notice the kid before I started the dive? If I did notice the kid, was my dive too far along for me to stop or change my direction if I had wanted to? Presumably, the answer to the first question will sometimes be 'no,' and the answer to the second will sometimes be 'yes.' In these cases it does indeed seem that my action is not the product of my free choice. On the other hand, things could be different. Suppose that I did notice what was going on, but made the dive anyway because I wanted to. Suppose I also approved of that desire (because this was my only chance for a gold medal!), and also that I would have stopped or changed direction if I had wanted to. In that case, it seems—at least arguably—that I *did* act of my own free will. But what sort of *consciousness*, if any, is required for me to have acted in this way?

Consider Frankfurt's Condition (i): I did what I wanted to do; that is, my dive was caused by my wanting to dive (or it was produced, at least, by a 'conditional intention' to dive under certain circumstances). Now if, as Chalmers contends, an unconscious creature is capable of the 'deliberate control of behavior,' then it seems that it could

9. Even Wu (2013a), who is sanguine about the possibility of agency without conscious awareness, acknowledges that *free* agency may require something more. See his (2013, 258): 'There is no denying we are often moved to act, and on the antecedents of action rest important questions about the rationality, morality, and freedom of our actions. But all these are higher-order properties of agency. Agency itself is an internal feature of certain processes....'

do what it wants in the same way that I did. Frankfurt's second condition, however, may seem harder to meet, since it involves second-order desires. That is, to act 'of one's own free will,' rather than merely 'acting freely,' requires that the agent want to act on her effective desires (or at least approve of so acting), and this requires a capacity for *reflecting on* the (1st order) desires she has. After all, for Frankfurt, non-human animals may be capable of free action in the sense that nothing prevents them from doing what they are inclined to do, but they are incapable of acting on their own free will because they do not possess the reflective capacities that would allow them to take a stand on the desires (or inclinations) they have. And surely, one might think, to meet this 'hierarchical' condition on freedom one must be *consciously aware* of one's desires.

But is the connection between having reflective capacities of this sort and being (phenomenally) conscious so clear? Consider once again Chalmers's imagined creatures: creatures that meet conditions (1)-(6), but with no (phenomenally) conscious experience. As mentioned above, such creatures are supposed to have the capacity not only to behave the way we do, but also (among other things) to 'discriminate, categorize, and react to environmental stimuli' and to access and report on their own mental states. It seems, therefore, that such creatures could have the reflective capacities required to meet all Frankfurt's conditions for having a will that is free. The same, it seems, can be said about views (e.g. Watson, 1975) that take free actions to be those compatible not with one's second-order desires, but with one's *values*, as long as values can be characterized in some naturalistic way.¹⁰

The same questions, moreover, can be raised for other contemporary characterizations of actions that arise from an agent's free will. Kadri Vihvelin, in her recent (2013) book, contends that an action arises from an agent's free will only if the agent could have done otherwise, where this requires having both the *ability* and *opportunity* to have done so. Vihvelin gives a subtle and rich characterization (176) of what counts as the sort of ability (or abilities) relevant to our concerns, namely, that one have 'an intrinsic disposition to do X in response to the stimulus of *one's trying to do X*.' Moreover, she continues, '[a]ll that's required for trying is that you acquire—somehow or other—an effective desire or intention; that is, a desire or intention that is causally effective [in certain specified ways].'¹¹ Here too, however, it's not clear that possessing (to coin a

10. Indeed, even outside the realm of Chalmersian beings, it is easier to think of situations in which people are blind to what they *really* consider to be important; self-deception abounds in the domain of value.

11. Once again, both Frankfurt and Vihvelin go on to argue that acts can meet these conditions in a deterministic universe, but I'm not interested in evaluating that contention here.

phrase) *Vihveibilities* requires some sort of (phenomenally) *conscious awareness*—either of one’s own motivations or the environmental conditions that afford opportunities to act. That is, it’s not clear why an unconscious creature that meets conditions (1)-(6) could not have these abilities too.

Yet another characterization of actions arising from free will is that the actions are *reasons-responsive*. Does this require phenomenal consciousness either of one’s own mental states, or items in one’s environment? Let’s look more closely at what it takes for an action to be reasons-responsive. One possibility is that the agent ‘acts for reasons’ in the sense that her action must be caused not just by a desire, but by a *rational* desire—that is, a desire that (either) coheres sufficiently with the rest of her desires (and beliefs and perhaps values) or with the desires, beliefs, and values that meet certain independent conditions of rationality. Either way, it’s not clear that an unconscious creature that meets conditions (1)-(6) could not act for reasons in this sense, even though it is completely (phenomenally) unconscious.

Another way to think about reasons-responsiveness is that an agent must be able to modify her actions if she *recognizes that there are reasons to do so*—whether these are beliefs and desires that she has suppressed or otherwise not yet noticed, or environmental conditions that turn out to be different than expected. But once again it’s unclear why an unconscious creature that meets conditions (1)-(6) couldn’t do *that*.

The upshot, thus, is that it seems that a creature without (phenomenally) conscious awareness could be capable of acting freely, or acting from its own free will, on any of these causal, structural, or counterfactual accounts of freedom.¹²

However, even if it’s possible for such creatures to exercise free will, we are not beings of this sort; we have conscious access both to our own experiences and the world around us, and it seems that many of our actions, including those studied by Libet, arise from a conscious wish or inclination. We have the experience of agency (Horgan, 2011) or of conscious will (Wegner, 2002). And if these conscious inclinations do not in fact initiate our actions, but our actions nonetheless count as arising from free will, then

12. Of course there are other views of what’s required for free will; for example, that the action be the product of *agent causation*; that is, that it be caused not by events such as beliefs, desires, or intentions, but by the agent herself. It’s a matter of debate whether the idea of agent causation is coherent; but it’s not clear why a creature of the sort we’ve been discussing couldn’t be the cause of its actions in just the way that a reasonable theory of agent causation demands. Granted, often people object to the possibility that a robot or similar machine could have free will—but this is usually on the grounds that such a creature would be acting according to the way it was programmed, and not because there are questions about whether such creatures could be phenomenally conscious.

our experiences of agency, or of the efficacy of our conscious will, must be regarded as illusory—just as Libet, Wegner, and many others have argued, and this would be disturbing in itself.

Perhaps this is so (though some, e.g. Paglieri, 2013, and Gallagher, 2013, have argued that the phenomenology of agency is not as robust as Libet and Wegner suggest).¹³ In any case, the illusoriness of the experience of agency, if it obtains (remember that there is dispute about whether Libet's data show that his subjects' actions do not arise from their conscious inclinations), is compatible with the *existence* of human free will. And if our experiences of agency are systematically erroneous, then they join a fairly large club. It's not unusual for us to be wrong about the causes of a variety our actions—and about other phenomena presented in introspection.

To be sure, it would be disturbing to think that our actions may be caused by some sort of nefarious manipulator while it seems to us that they are caused by our conscious decisions. However, what makes this possibility so disturbing, I suggest, is not that that the causes of our actions are not conscious, but that they are motivations that we would not want to be effective, or that diverge from our values, or are inconsistent with the motivations we identify with, or take to reflect our real or 'deep' selves. But, according to the best-known views about what makes an action free, these would not be cases in which we err about the motivations for our free actions, but rather cases in which our actions are not in fact free.

However, I suspect that things may seem different for attributions of *moral responsibility*. It's harder, it seems, to blame or praise people for what they do if they have no (phenomenally) conscious experience either of their desires and inclinations (be they conflicting or consistent) or of the environment that is prompting them, given those desires and inclinations, to act in certain ways.

Consider the sorts of things we say to excuse people from blame (or question whether they are being legitimately praised): 'She wasn't aware of what she was doing'; 'He didn't recognize the difference between right and wrong.' These excuses, it seems, carry with them *not only* the suggestion that the agent wasn't responsive, in some way, to environmental (and internal) exigencies, but also that the agent did not have

13. As Paglieri puts it (2013, 136), 'There is no proof that free actions are phenomenologically marked by some specific 'freedom attribute...and thus this non-existent entity cannot be invoked to justify our judgments on free will and agency, be they correct or mistaken.' He goes on to suggest that we regard our acts as free by default, and only consider that they are not free if we have countervailing evidence, such as the experience of coercion. Gallagher (2013) argues that the sense of agency can sometimes arise after retrospective reflection, and often has a social dimension.

a robust conscious awareness of them, the sort that comes with the ability to adopt another person's subjective situation, or point of view. To be able to do this, of course, requires that one *have* a point of view, in the sense of having phenomenally conscious experiences of (and, perhaps more important, emotional responses to) the things that are the causes and effects of one's actions. If so, then moral agency may require (phenomenal) consciousness. That is, it may seem that unconscious creatures (with the relevant functional capacities) may be able to be intentional agents, indeed free agents—but not moral agents.¹⁴

This view has at least some support in the literature on moral responsibility. For example, T.M. Scanlon (1998, 281) argues that although a computer could be regarded as responsible 'in a causal sense for the processes it governs'...we would not 'regard it as "responsible" in the sense responsible for moral blame.' And this, he continues (282) is 'because computers, even very sophisticated ones, ...are not conscious.'¹⁵

If acts produced by unconscious motivations could nonetheless be free, however, this would not be the only case in which actions could be regarded as free but not blameworthy (or praiseworthy). Think, for example, of actions that are coerced: Someone holds a gun to your head and says 'Your money or your life'—and you hand over your money. Although some claim that you don't do so freely, it is equally intuitive to hold

14. In addition, there are views that contend that having certain *emotional* responses is necessary for recognizing others as bona-fide moral agents who can legitimately be blamed and praised for what they do. See P.F. Strawson (1962).

15. Indeed, Scanlon goes further and argues (282) that 'it is crucial to a creature's being a rational creature that conscious judgment is one factor affecting its behavior. Computers, even very sophisticated ones, do not strike us as moral agents or rational creatures, partly because we believe that they are not conscious at all—and that there is no such thing as how reasons, or any other things, seem to them.' (I am grateful to Pamela Hieronymi for calling these passages to my attention.) Scanlon's mention of how things *seem* to an agent suggests that he means by 'consciousness' just what Chalmers takes to be independent of the capacities (1)-(6) discussed in the text. On the other hand, Michael S. Moore (2011, 223, col. 1) argues that even in the absence of 'phenomenal' (that is, Chalmersian) consciousness, the possession of 'dispositional' consciousness of one's motivation—that is, 'the ability to direct attention and to state that of which one is conscious, abilities that seem included in Chalmers's (1)-(6)—may be sufficient for moral responsibility. Nonetheless, Moore (223, col. 2) argues (with respect to Freudian explanation) that 'Although there may be truly unconscious agency that is nonetheless the agency of a person, that person's responsibility is not increased by virtue of such truly unconscious actions, intentions, or tryings.' It's not clear, however, just what counts as 'truly unconscious agency,' and whether it could be possessed by a creature that satisfies any of the Frankfurtian or other well-known conditions for the possession of free will. In any case, Moore goes on to argue that, in the Libet cases, agents can be regarded as acting because of willing that is conscious in both the phenomenal and dispositional sense.

that you *do* make a free choice and thereby act freely, but since the costs and benefits of your available alternatives have been manipulated in certain ways, you shouldn't be held responsible for what you do (or choose). Your actions could be 'yours' or 'up to you' metaphysically, but not in ways that make a difference to praise, blame, or other sorts of moral evaluation.

If moral responsibility does not follow automatically from freedom in cases of coercion, then perhaps it does not follow from freedom in Chalmersian (or blindsight) cases, either. Indeed, the separation of questions about free will from questions about moral responsibility may make certain discussions of Libet's results more intelligible. For example, in his (2013b), Mele considers a situation in which an agent consciously deliberates among alternatives but makes no decision about them, and takes up the deliberation days later, still feeling 'unsettled' about what to do. Mele suggests that if the agent were to find out that she *unconsciously decided to do A after her initial period of deliberation*, then she may doubt that she did A freely. He argues, however, that this case is significantly different from a case in which the time lag between the action and the prior decision (made after conscious deliberation) is much shorter, on the order of half a second. In this case, Mele argues, the time lag is not threatening, since we can think that our detection of the decision merely 'lags a bit behind the actual decisions.' He acknowledges that it may be odd to hang freedom of action on the time interval between decision and action in this way, and suggests (786) that the presence of a long time lag between decision and awareness prompts skepticism about whether conscious deliberation did in fact play a significant role in the production of the decision. It may make more sense, however, to argue that the time lag makes a difference not to the determination of whether the action was free, but to the attribution of *moral responsibility* to the agent.¹⁶

Now I myself do not have firm views about whether moral responsibility requires conscious awareness of one's motives (or anything else); I'm willing to consider the possibility that unconscious agents (who meet the conditions for freedom) are morally responsible as well. But perhaps the recognition that there may be a gap between freedom and responsibility will make it easier to get a clear-eyed account of what is required for freedom—and, perhaps more important, a clear-eyed account of the aims of praise and punishment. If we separate considerations of whether an individual is conscious from questions about whether she is free, then we may be able to get a better

16. Moreover, it can help to make sense of Wegner's claim, highlighted by Dennett (2003, 242) that '[i]llusory or not, conscious will is the person's guide to his or her own moral responsibility for action.'

hold not just on questions about freedom, but also on questions of moral responsibility, and therefore, this may be a view that is worthy of further discussion.¹⁷

17. The idea that we should firmly separate considerations of the aims of praise and punishment from consideration of whether an agent acts freely is familiar from the work of hard determinists, such as Holbach, who argues that this is required if people *cannot* exercise free will. It may be required as well for those who hold that humans can act freely in a wider variety of circumstances than we had initially thought. In addition, the separation of considerations about whether an individual is conscious from considerations about whether she is free—or could even be an agent—may have salutary effects for further theorizing about other phenomena that seem to be tied to action or intentional agency, such as attention.

References

- Block, Ned. 1995. "On a Confusion about the Function of Consciousness." *Behavioral and Brain Sciences* 18 (2): 227–87.
- Chalmers, David. 1995. "Facing up to the Problem of Consciousness." *Journal of Consciousness Studies* 2 (3): 200–219.
- Clark, Andy, Julian Kiverstein, and Tillmann Vierkant. 2013. *Decomposing the Will*. New York: Oxford University Press.
- Dennett, Daniel C. 2003. *Freedom Evolves*. New York: Penguin Group (Viking).
- Frankfurt, Harry. 1971. "Freedom of the Will and the Concept of a Person." *Journal of Philosophy* 68 (1): 5–20.
- Gallagher, Shaun. 2013. "Ambiguity in the Sense of Agency." In *Decomposing the Will*, edited by Andy Clark, Julian Kiverstein, and Tillmann Vierkant, 118–135. New York: Oxford University Press.
- Horgan, Terry. 2011. "The Phenomenology of Agency and the Libet Results." In *Conscious Will and Responsibility: A Tribute to Benjamin Libet*, edited by Walter Sinnott-Armstrong and Lynn Nadel, 159–172. New York: Oxford University Press.
- Libet, Benjamin. 1985. "Unconscious cerebral initiative and the role of conscious will in voluntary action." *Behavioral and Brain Sciences* 8 (4): 529–66.
- Libet, Benjamin. 2011. "Do We Have Free Will?" In *Conscious Will and Responsibility: A Tribute to Benjamin Libet*, edited by Walter Sinnott-Armstrong and Lynn Nadel, 1–10. New York: Oxford University Press.
- Mele, Alfred. 2011. "Libet on Free Will: Readiness Potentials, Decisions, and Awareness." In *Conscious Will and Responsibility: A Tribute to Benjamin Libet*, edited by Walter Sinnott-Armstrong and Lynn Nadel, 23–33. New York: Oxford University Press.
- Mele, Alfred. 2013a. "Free Will and Neuroscience." *Philosophical Exchange* 43 (1): Article 3.
- Mele, Alfred. 2013b. "Unconscious Decisions and Free Will." *Philosophical Psychology* 26 (6): 777–789.
- Mele, Alfred. 2013c. "Vetoing and Consciousness." In *Decomposing the Will*, edited by Andy Clark, Julian Kiverstein, and Tillmann Vierkant, 73–86. New York: Oxford University Press.
- Mele, Alfred. 2014. *Free: Why Science Hasn't Disproved Free Will*. New York: Oxford University Press.

- Nadelhoffer, Thomas. 2011. "The Threat of Shrinking Agency and Free Will Disillusionism." In *Conscious Will and Responsibility: A Tribute to Benjamin Libet*, edited by Walter Sinnott-Armstrong and Lynn Nadel, 173–188. New York: Oxford University Press.
- Paglieri, Fabio. 2013. "There's Nothing Like Being Free: Default Dispositions, Judgments of Freedom, and the Phenomenology of Coercion." In *Decomposing the Will*, edited by Andy Clark, Julian Kiverstein, and Tillmann Vierkant, 136–159. New York: Oxford University Press.
- Scanlon, T.M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Sinnott-Armstrong, Walter, and Lynn Nadel. 2011. *Conscious Will and Responsibility: A Tribute to Benjamin Libet*. New York: Oxford University Press.
- Strawson, Galen. 2003. "Mental ballistics or the involuntariness of spontaneity." *Proceedings of the Aristotelian Society* 103: 227–256.
- Strawson, P.F. 1962. "Freedom and Resentment." *Proceedings of the British Academy* 48: 1–25.
- Vihvelin, Kadri. 2013. *Causes, Laws, and Free Will: Why Determinism Doesn't Matter*. New York: Oxford University Press.
- Wegner, Daniel M. 2002. *The Illusion of Conscious Will*. Cambridge, MA: MIT Press.
- Wu, Wayne. 2011. "Attention as Selection for Action." In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and and Wayne Wu, 97–116. New York: Oxford University Press.
- Wu, Wayne. 2013a. "Mental Action and the Threat of Automaticity." In *Decomposing the Will*, edited by Andy Clark, Julian Kiverstein, and Tillmann Vierkant, 244–261. New York: Oxford University Press.
- Wu, Wayne. 2013b. "The Case for Zombie Agency." *Mind* 122 (485): 217–230.