

PROXIMAL INTENTIONS, INTENTION-REPORTS, AND VETOING

1. *Introduction*

Benjamin Libet has argued that although free will does not initiate actions, it may be involved in “vetoing” conscious intentions or urges to act (1985, 1999, 2004, pp. 137-49). In this connection, Libet attempts to generate evidence about when his subjects become aware or conscious of pertinent intentions or urges. His method is to instruct subjects to perform a flexing action whenever they wish while watching a rapidly revolving dot on a clock face and to report later – after they flex – on where the dot was when they first became aware of their intention or urge to flex (1985). Libet found (1985, p. 532) that the average time of reported initial awareness was 200 milliseconds (msec) before the time at which an electromyogram (EMG) shows relevant muscular motion to begin (time 0). The question how accurate the subjects’ reports are likely to have been has received considerable attention (for a review, see van de Grind 2002). It is the topic of section 5.

The intentions and urges that concern Libet are *proximal* intentions and urges (Mele, 1992) – that is, intentions or urges to do things *at once*. The following labels will facilitate discussion:

I-time: The time of the onset of a proximal *intention* to flex.

A-time: The time of the onset of the subject’s *awareness* of such an intention.

B-time: The time the subject *believes* to be A-time when responding to the experimenter’s question about A-time.

How are these times related? Libet’s position is that average I-time is 550 msec before time 0 (that is, -550 msec) for subjects who are regularly encouraged to flex spontaneously and who report no “preplanning” of their movements, average A-time is -150 msec, and average B-time is -200 msec (1985, p. 532, 2004, pp. 123-26).¹ However, researchers who treat intentions as being, by definition, conscious states identify I-time with A-time. And some who identify these times have suggested that A-time is too late to permit “motor intentions” to be among the causes of actions (Lau *et al.*, 2007, p. 81). This issue about motor intentions and A-time is the topic of section 3.

Whether subjects have time to veto conscious proximal intentions or urges, as Libet claims, is the topic of section 4. Section 2 describes a recent study that bears on all three of the issues raised here – the accuracy of subjects’ reports, the connection between motor intentions and actions, and the question about vetoing.

2. *A Recent Study*

A recent study by Hakwan Lau, Robert Rogers, & Richard Passingham is motivated partly by a reference (2007, p. 81) to the following comment by Daniel Wegner on Libet’s results: “The position of conscious will in the time line suggests perhaps that the experience of will is a link in a causal chain leading to action, but in fact it might not even be that. It might just be a loose end – one of those things, like the action, that is caused by prior brain and mental events” (2002, p. 55). Lau *et al.* observe that Wegner “does not show that motor intentions are in fact *not* causing the actions” and that “if intentions, in fact, arise after the actions, they could not, in principle, be causing the actions” (p. 81).

The main experiment (Experiment 1) reported in Lau *et al.* 2007 combines Libet's "clock paradigm" with the application of transcranial magnetic stimulation (TMS) over the presupplementary motor area. The dot on a Libet clock revolves at 2560 msec per cycle. While watching such a clock, subjects pressed a computer mouse button "at a random time point of their own choice" (p. 82). In the "intention condition," after a delay of a few seconds, subjects were required to move a cursor to where they believed the dot was "when they first felt their intention to press the button." In the "movement condition," they followed the same procedure to indicate where they believed the dot was "when they actually pressed the button." There were a total of 240 trials per subject. TMS was applied in half of the trials. Half of the applications occurred "immediately after action execution" and half occurred at a delay of 200 msec. There were ten subjects.

The results were as follows. When TMS was not applied, "the group mean for the judged onset of intention was -148 msec . . . relative to the time of the recorded button press" and the mean for the movement judgment was -50 msec (Lau *et al.*, 2007, p. 83).² "The effect of TMS for each individual was assessed by subtracting the median judgment value for the [non]TMS trials from that of the TMS trials." In the intention condition, "the group mean for the TMS effect was -9 msec" when TMS was not delayed and -16 msec when TMS was delayed by 200 msec. In the movement condition, the mean for the effect was 14 msec when TMS was not delayed and 9 msec when TMS was delayed by 200 msec.

Lau *et al.* conducted a second experiment with ten subjects "to test whether the effect obtained in Experiment 1 was actually due to memory or responding, rather than the experienced onset itself" (2007, p. 84). In this experiment, application of TMS occurred either 500 msec after the button press or between 3280 and 4560 msec after it. When TMS was not applied, the group mean for the judged onset of the feeling of intention was -110 msec and the mean for the movement judgment was -7. In the intention condition, "the group mean for the TMS effect was 9 msec" for the 500 msec delay and 0 msec for the longer delay. In the movement condition, the mean effect was 5 msec for both delays. As Lau *et al.* report, "the effect observed in Experiment 1, that is, the exaggeration of the difference of the judgments for the onsets of intention and movement," was not observed in Experiment 2 (p. 87).³

As Lau *et al.* view matters, "the main question is about whether the experience of intention is fully determined before action execution" (2007, p. 87). Their answer is no: "The data suggest that the perceived onset of intention depends at least in part on neural activity that takes place after the execution of action" (p. 89).

3. *Motor Intentions and Actions*

One question raised in my introduction is whether motor intentions to press a mouse button, for example, are among the causes of the pressing actions. Lau *et al.* take their results to bear on this question. As I mentioned, they point out that Wegner has not shown that "motor intentions are in fact *not* causing the actions" and they observe that "if intentions arise after actions, they could not, in principle, be causing the actions" (2007, p. 81). Lau *et al.* seemingly believe that they have produced evidence that their subjects' proximal intentions to press a mouse button arise after they press. But have they?

Return to the notions of I-time and A-time:

I-time: The time of the onset of a proximal intention to flex.

A-time: The time of the onset of the subject's awareness of such an intention.

The time of the *onset* of x , for any x , is the time of x 's initial presence. Onsets of x are, of course, to be distinguished from *causes* of onsets of x . Libet contends that I-time precedes A-time by over a third of a second. In some of his studies, subjects are instructed to flex their right wrists whenever they wish and electrical readings from the scalp (EEGs) – averaged over at least 40 flexings for each subject – show a shift in readiness potentials (RPs) beginning at about -550 msec (1985, pp. 529-30). (These RPs are called “type II RPs” [p. 531].) Subjects are also instructed to “recall . . . the spatial clock position of a revolving spot at the time of [their] initial awareness” (p. 529) of something, x , that Libet variously describes as a “decision,” “intention,” “urge,” “wanting,” “will,” or “wish” to move.⁴ On average, “RP onset” preceded what the subjects reported to be the time of their initial awareness of x by 350 msec. Libet maintains that a proximal intention to flex emerges in the brain at -550 msec, well before the subject becomes conscious of it (p. 536). The “intention to act arises unconsciously” (p. 539).

I am not claiming that Libet is correct to identify what is signified by the beginning of the type II RP as the onset on a proximal intention to flex. In fact, I have argued elsewhere, on the basis of additional data, that Libet's claim is implausible (Mele, 2006, ch. 2). Even so, it should be noted that *if* I-time precedes A-time, then *even if* A-time is “after actions” (Lau *et al.*, 2007, p. 81), I-time might be early enough for motor intentions to play a role in causing actions. Possibly, Lau *et al.* are assuming with Wegner (2002, p. 18) that nonconscious items cannot be intentions. But this theoretical assumption has been challenged (Marcel, 2003, Mele, 2004). Researchers who think of proximal intentions in terms of their functional roles tend to see no need for agents to be conscious of all of their proximal intentions (Mele, in press). The last time you signaled for a turn in your car, were you conscious of an intention to do that? Probably not. But in the absence of a stipulative definition of intention that requires intenders to be conscious of all of their intentions, it does not follow from your not being conscious of an intention to signal that you had no such intention.

Some neuroscientists are doing research with monkeys that is aimed at developing neural prostheses that will enable paralyzed people to move their limbs. They seek to record monkeys' *intentions* and to use the associated neural signals to move robot arms, for example (see Andersen & Buneo, 2002, Musallam *et al.*, 2004). Now, in the words of Richard Passingham & Hakwan Lau, “the operational index of consciousness is the ability to report” (2006, p. 67). The monkeys in these experiments report nothing, of course. Yet, in the articles I cited, the experimenters express no reservations about attributing intentions to them. Presumably, if they had thought that all intentions are conscious intentions, they would have worried about the absence of evidence of the required consciousness.

Just as your nose is one thing and your awareness of it is another, might not a subject's proximal intention to press a button be one thing and his or her awareness of that intention be another? If so, two points need to be noted. First, the onset of our awareness of many of our proximal intentions – in those situations in which we are aware of them – may be significantly later than the onset of the intentions. Second, our proximal intentions may play causal roles that our awareness of them does not play.

Again, Lau *et al.* contend that “The data suggest that the perceived onset of intention depends at least in part on neural activity that takes place after the execution of action” (2007, p. 89). But it does not follow from this that the data suggest that the *actual onset* of intention depends at least partly on neural activity that takes place after the execution of action. Even if Lau *et al.* are right about what the data suggest, the data leave it wide open that the actual onset of the intention is early enough for the intention to be among the causes of the action. Moreover, as I will explain later, their data also leave it wide open that the time of the onset of subjects’ *awareness* of proximal intentions (A-time) does not depend at all on “neural activity that takes place after the execution of the action” and that A-time sometimes is early enough for motor intentions of which agents are conscious to play a role in causing actions.

4. *Vetoing*

What about awareness of proximal intentions (or urges) to flex or press a button in the studies by Libet and Lau *et al.*? When might that emerge? And does it ever emerge early enough for subjects to veto conscious proximal intentions (or urges) to flex, as Libet claims? (Libet associates veto power with some pretty fancy metaphysics: see, for example, Libet 1999. I set the metaphysical issues entirely aside here.)

To veto a conscious proximal intention or urge to do something is to decide not to act on it and to refrain, accordingly, from acting on it. The question here is whether there ever is enough time for this in the experimental contexts at issue. In this connection, one should ask exactly what Lau *et al.* mean by their suggestion that “the perceived onset of intention depends at least in part on neural activity that takes place after the execution of action” (2007, p. 89). For example, do they mean to exclude the following: the subjects are aware of a proximal intention before they press the button even though they do not have a definite opinion about when they perceived the onset of their intention – or when they first felt the intention (see n. 2) – until after they act? Apparently not, for they grant that “it could be the case that some weaker form of experience of intention is sufficiently determined by neural activity that takes place before the execution of the action” (p. 89). If, in the experiments at issue, subjects (sometimes or always) are aware of proximal intentions (or urges) to press before they press, are they sometimes aware of this far enough in advance of the pressing to veto the conscious intention (or urge)?

Someone who wishes to answer the preceding question may seek evidence about the time of the onset of awareness of the intention (or urge) – A-time – and evidence about how much time one needs to veto a proximal intention or urge. Libet says that his subjects were “free *not* to act out any given urge or initial decision to act; and each subject indeed reported frequent instances of such aborted intentions” (1985, p. 530). Might these subjects have been right – at least sometimes? Again, Libet claims that average A-time in his subjects is -150 msec. The plausibility of that claim is one of the topics of the following section. Libet believes that a subject who becomes aware of a proximal intention or urge to flex at -150 msec has enough time for a veto (1985, pp. 537-38). What evidence do we have that he is right or wrong about this, and how might we get additional evidence?

It is advisable to distinguish between the following two claims:

1. An A-time of -150 msec leaves enough time for a veto.
2. Libet's subjects sometimes have enough time for a veto.

Even if claim 1 is false, claim 2 might be true, if the A-time for Libet's subjects is sometimes earlier than -150 msec. One piece of evidence that Libet offers for the truth of claim 2 is subjects' reports that they frequently vetoed proximal intentions or urges. Naturally, researchers who believe that there is not sufficient time for vetoing will be skeptical about these reports.

Libet also claims to find support for the truth of claim 2 in a veto study. Subjects are instructed to prepare to flex their fingers at a prearranged time (as indicated by a revolving dot on a Libet clock) and "to veto the developing intention/preparation to act . . . about 100 to 200 ms before the prearranged clock time" (1985, p. 538). Subjects receive both instructions at the same time. Libet writes:

a ramplike pre-event potential was still recorded . . . resembl[ing] the RP of self-initiated acts when preplanning is present. . . . The form of the 'veto' RP differed (in most but not all cases) from those 'preset' RPs that were followed by actual movements [in another experiment]; the main negative potential tended to alter in direction (flattening or reversing) at about 150-250 ms before the preset time. . . . This difference suggests that the conscious veto interfered with the final development of RP processes leading to action. . . . The preparatory cerebral processes associated with an RP can and do develop even when intended motor action is vetoed at approximately the time that conscious intention would normally appear before a voluntary act. (1985, p. 538)⁵

Keep in mind that the subjects were instructed in advance *not* to flex their fingers, but to prepare to flex them at the prearranged time and to "veto" this. The subjects intentionally complied with the request. They intended from the beginning *not* to flex their fingers at the appointed time. So what is indicated by the segment of what Libet refers to as "the 'veto' RP" that precedes the change of direction? Presumably, not the presence of an *intention* to flex; for then, at some point in time, the subjects would have both an intention to flex at the prearranged time and an intention not to flex at that time. And how can a normal agent simultaneously intend to *A* at *t* and intend not to *A* at *t*?⁶ In short, it is very plausible that Libet is mistaken in describing what is vetoed as "*intended motor action*" (p. 538; my emphasis). And if he is mistaken about that, he has not provided evidence about how much time is enough time for the vetoing of a proximal *intention*.

In some talks I have given on Libet's work, I tell the audience that I will count from 1 to 5, and I ask them to prepare to snap their fingers when I say "5" but not to snap them. (When I hear no finger snapping, I jokingly praise my audience for being in control of their fingers.) Someone might suggest that these people have *conscious intentions not to flex* when I get to 5 and *unconscious intentions to flex* then and that the former intentions win out over the latter. But this suggestion is simply a conjecture – an unparsimonious one – that is not backed by evidence. We have every reason to believe that the members of my audience do *not* intend to snap their fingers and, similarly, that the subjects in Libet's veto experiment do *not* intend to flex. Accordingly, we have good reason to refuse to grant that the veto study provides information about how many

milliseconds are enough to permit the vetoing of a proximal *intention*.

If we were to think of a decision not to act on a conscious proximal urge as an analogue of a stop signal and the onset of awareness of the urge as an analogue of a go signal, we might seek indirect evidence about how much time is needed for the vetoing of a proximal urge from stop-signal studies. Imagine a study in which one tone signals subjects to click a mouse button straightaway and another tone signals them not to click. How long after the former tone sounds can the latter tone sound and be effective? Gordon Logan reports that “eye movements, hand movements, key presses, squeezes, and speech can all be stopped in about 200 msec” (1994, p. 191). *If* conscious proximal urges or intentions to flex a wrist or to press a mouse button can be vetoed in about 200 msec, are subjects in Libet’s or Lau’s studies ever aware of relevant proximal urges or intentions early enough to veto them? Obviously, that depends on the time of the onset of the subjects’ awareness of the proximal urges or intentions, an issue to which I return in the following section.

Marcel Brass & Patrick Haggard conducted an experiment in which subjects “were instructed to freely decide when to execute a key press while observing a rotating clock hand” on a Libet-like clock and “to cancel the intended response at the last possible moment in some trials that they freely selected” (2007, pp. 9141-42). They report that “the mean proportion of inhibition trials was 45.5%, but that there were large interindividual differences, with the proportion of inhibition trials ranging from 28 to 62%” and that “subjects reported the subjective experience of deciding to initiate action a mean of -141 ms before the key press on action trials” (p. 9142). If the subjects actually did what Brass & Haggard say they were instructed to do, they vetoed their decisions an average of 45.5% of the time.

In light of Brass & Haggard’s results, should everyone now grant that Libet was right – that people have time to veto conscious proximal decisions or intentions? Naturally, some researchers will worry that, “in inhibition trials,” subjects were *simulating* vetoing conscious proximal decisions rather than actually making conscious proximal decisions to press that they proceeded to veto. A reasonable question to ask in this connection is what strategy subjects thought they were adopting for complying with the instructions. There are various possibilities, and four of nineteen subjects in a “pre-experiment” were excluded from the actual experiment because they “reported that they were not able to follow the instructions” (2007, p. 9141). Apparently, these subjects failed to hit upon a strategy that they deemed satisfactory for complying with the instructions. What strategies might the other fifteen subjects have used?

Here is one hypothetical strategy:

Strategy 1. On each trial, consciously decide in advance to *prepare* to press the key when the dot hits a certain point p , but leave it open whether, when the dot hits p , I will consciously decide to press right then or consciously decide not to press on that trial. On some trials, when the dot hits p , decide right then to press at once; and on some other trials decide right then not to press. Pick different p points on different trials.⁷

Subjects who execute this strategy as planned do not actually veto conscious proximal decisions to press. In fact, they do not veto any conscious decisions. Their first conscious

decision on each trial is to *prepare* to flex a bit later, when the dot hits point *p*. They do not veto this decision; they do *prepare* to flex at that time. Nor do they veto a subsequent conscious decision. If, when they think the dot reaches *p*, they consciously decide to press, they press; and if, at that point, they consciously decide not to press, they do not press.

A second strategy is more streamlined:

Strategy 2. On some trials, consciously decide to press the key and then execute that decision at once; and on some trials, consciously decide not to press the key and do not press it.

Obviously, subjects who execute this strategy as planned do not veto any conscious decisions.

Here is a third strategy:

Strategy 3. On some trials, consciously decide to press the key a bit later and execute that decision. On other trials, consciously decide to press the key a bit later but do not execute that decision; instead veto (cancel, retract) the decision.

Any subjects who execute this strategy as planned do veto some conscious decisions; but the decisions they veto are not proximal decisions. Instead, they are decisions to press *a bit later*. A subject may define “a bit later” in terms of some preselected point on the clock or leave the notion vague.

The final hypothetical strategy to be considered is even more ambitious:

Strategy 4. On some trials, consciously decide to “press now” and execute that decision at once. On other trials, consciously decide to “press now” but do not execute that decision; instead immediately veto (cancel, retract) the decision.

If any subjects execute the fourth strategy as planned, they do veto some conscious proximal decisions. But, of course, we are faced with the question whether this strategy is actually executable. Do we have enough time to prevent ourselves from executing a conscious proximal decision to press?⁸

The results of Brass & Haggard’s experiment leave this question unanswered. If we knew that some subjects were successfully using strategy 4, we would have an answer. But what would knowing that require? Possibly, if asked about their strategy during debriefing, some subjects would describe it as I have described strategy 4. But that alone would not give us the knowledge at issue. People have been wrong before about how they do things.

5. Accuracy

How accurate are subjects’ reports about when they first became aware of a proximal intention or urge likely to have been? Framed in terms of A-time (the time of the onset of the subject’s awareness of a proximal intention) and B-time (the time the subject believes to be A-time when answering the experimenter’s question about A-time), the question about intentions is this: How closely does B-time approximate A-time?

When, according to Lau *et al.*, do their subjects first become aware of (or “perceive” or “feel”) “their [proximal] intention to press the button” (2007, p. 82). That is, what is A-time according to Lau *et al.*? They suggest that “the perceived onset of intention depends on neural activity that can be manipulated by TMS as late as 200 msec after the execution of a spontaneous action” (p. 89). But this suggestion does not itself amount to the suggestion that awareness of an intention arises only after the action. Again, Lau *et al.* grant that “it could be the case that some weaker form of experience of intention is sufficiently determined by neural activity that takes place before the execution of the action.” As I mentioned, it is possible that subjects are aware of a proximal intention before they press the button even if they do not have a definite opinion about when they first felt their intention until after they act. Lau *et al.* do not offer a definite answer about what average A-time is in their subjects; and they are right not to do so, because the evidence they produce does not warrant such an answer. From the fact that B-time can be manipulated by TMS after an action, nothing follows about what average A-time is. And because this is so, the former fact also leaves it open that A-time sometimes is early enough for proximal intentions of which subjects are conscious to be among the causes of their actions.

It may be that the beliefs that subjects report in response to the question about A-time are always a product of their experience of a proximal intention (or decision or urge), their perception of the clock around the time of the experience just mentioned, and subsequent events. They may always estimate A-time based on various experiences rather than simply “remembering” it. And making the estimate is not a particularly easy task. Reading the position of a rapidly revolving dot at a given time is no mean feat, as Wim van de Grind observes (2002, p. 251). The same is true of relating the position of the dot to such an event as the onset of one’s awareness of a proximal intention to press a button.

Perhaps the difficulty of the task helps to account for the considerable variability in B-times that individuals display across trials. Patrick Haggard & Martin Eimer (1999) provide some relevant data. For each of their eight subjects, they located the median B-time and then calculated the mean of the premedian (i.e., “early”) B-times and the mean of the postmedian (i.e., “late”) B-times. At the low end of variability by this measure, one subject had mean early and late B-times of -231 and -80 and another had means of -542 and -351 (p. 132). At the high end of variability, one subject’s figures were -940 and -4 and another’s were -984 and -253. Bear in mind that these figures are for means, not for extremes. These results do not inspire confidence that B-time closely approximates A-time. If we had good reason to believe that A-times – again, times of actual onsets of a subject’s awareness of a proximal intention – varied enormously across trials for the same subject, we might not find enormous variability in a subject’s B-times worrisome in this connection. But we do not have good reason to believe this. Moreover, Lau *et al.* have produced evidence that B-times can be affected by neural activity that occurs “after the execution of a spontaneous action” (2007, p. 89); and, of course, no A-time that precedes “the execution of a spontaneous action” can be affected by anything that happens after the action. (Whether it ever happens that subjects become aware of a proximal intention to do something only after they do it is an open question. But that question is not answered by subjects’ reports of B-times.)

One naturally wonders whether subjects who display relatively low variability in their B-times use different strategies for complying with the instructions than subjects

who display relatively high variability. In future studies, asking subjects what they took their strategies to be might prove informative. Haggard notes that “the large number of biases inherent in cross-modal synchronization tasks means that the perceived time of a stimulus may differ dramatically from its actual onset time. There is every reason to believe that purely internal events, such as conscious intentions, are at least as subject to this bias as perceptions of external events” (2006, p. 82). Perhaps some subjects’ strategies generate less bias than other subjects’ strategies.

Is there a way to make the subjects’ task a bit easier while also moving them a bit closer to something that might be regarded as simply remembering where the dot was at the onset of awareness of some pertinent mental event? Consider the following instruction set:

One way to think of deciding to press the button now is as consciously saying “now!” to yourself silently in order to command yourself to press the button at once. Consciously say “now!” silently to yourself whenever you feel like it and then immediately press the button. Look at the clock and try to determine as closely as possible where the dot is when you say “now!” Don’t count on yourself simply to remember where the dot is then. Instead, actively make a mental note of where the dot is when you say “now!” and try to keep that note in mind. You’ll report that location to us after you press the button.

A-time is not directly measured. Instead, subjects are asked to report, after the fact, what they believe A-time was. This is a report of B-time. Is there a way of gathering evidence about which of various alternative sets of instructions might help to yield B-times that are more reliable indicators of A-times? Some additional background sets the stage for an answer.

As Haggard observes, subjects’ reports about their intentions “are easily mediated by cognitive strategies, by the subjects’ understanding of the experimental situation, and by their folk psychological beliefs about intentions” (2006, p. 81). He also remarks that “the conscious experience of intending is quite thin and evasive” (2005, p. 291). Even if the latter claim is an overstatement and some conscious experiences of intending are robust, the claim may be true of many of the experiences at issue in Libet-style studies. One can well imagine subjects in Libet’s or Lau *et al.*’s experiments wondering occasionally whether, for example, what they are experiencing is an urge to act or merely a thought about when to act or an anticipation of acting soon. Again, Lau *et al.* say that they require their subjects to move a cursor to where they believed the dot was “when they first felt their *intention* to press the button” (2007, p. 82; emphasis mine). One should not be surprised if subjects given such an instruction were occasionally to wonder whether they were experiencing an intention to press or just an *urge* to press, for example. Subjects may also wonder occasionally whether they are actually *feeling* an intention to press or are mistakenly thinking that they feel such an intention. There is much less room for confusion and doubt about whether one is consciously saying “now!” to oneself. These observations generate a prediction. Subjects asked to report on when they said “now!” will – individually, but not necessarily collectively – exhibit significantly less variability in their reports (relative to time 0) than subjects asked to report on onsets of awareness of such things as intentions and urges. If the prediction is

shown to be true, we would have some reason to believe that their reports about when they consciously said “now!” involve less *guesswork*.

What about the “memory” part of the instruction set? It is difficult to predict what effect it would have. If the way people actually arrive at B-times is by estimating, after the fact, the pertinent position of the dot on the clock on the basis of various experiences, the “memory” instructions might be difficult to follow and result in greater variability. Fortunately, there is a way to find out: namely, by running the experiment. One set of experiments can use the “now!”-saying instructions in place of the “intention” instructions, and another can combine the “now!”-saying instructions with the “memory” instructions, as above. It would be interesting to see how the results of versions of Lau *et al.*’s Experiments 1 and 2 that use instructions like these compare to the results they reported.

I opened this section with the following question: How accurate are subjects’ reports about when they first became aware of a proximal intention or urge likely to have been? *Not very likely* certainly seems to be a safe answer. But there may be ways to improve accuracy.

6. Conclusions

My conclusions may be economically formulated in terms of I-time, A-time, and B-time.

I-time: The time of the onset of a proximal intention to flex.

A-time: The time of the onset of the subject’s awareness of such an intention.

B-time: The time the subject believes to be A-time when responding to the experimenter’s question about A-time.

Lau *et al.* have provided evidence that subjects’ B-times are influenced by the application of TMS at the time of a button press and 200 msec later. However this effect does not provide information about what I-time is nor about what A-time is. Consequently, their evidence about B-times does not support any of the following theses: proximal intentions always arise too late to be among the causes of actions; awareness of proximal intentions always arises too late for proximal intentions of which agents are conscious to be among the causes of actions; people never become aware of proximal urges or intentions early enough to veto them. Lau *et al.* correctly assert that Wegner “does not show that motor intentions are in fact *not* causing the actions” (p. 81). And they do not show this either – even regarding conscious proximal intentions. Even so, Lau *et al.*’s results are interesting. These results suggest that reports of B-times are reports of estimates that are based at least partly on experiences that follow action. Further study of this issue may shed light on the issue of the accuracy of subjects’ reports of A-time.

Brass & Haggard, for reasons set out in section 4, have not shown that their subjects actually veto conscious proximal decisions. Nor, for reasons set out in the same section, has Libet shown that his subjects have time to veto conscious proximal decisions, intentions, or urges. Whether subjects actually have time for such vetoing in the experimental settings at issue is still an open question.

Whether motor intentions are sometimes among the causes of some actions and whether there is enough time for people to veto conscious proximal urges, intentions, and decisions are interesting and important questions that have attracted the attention of

scientists and philosophers alike. In future work of the sort reviewed here on these questions, researchers would do well to bear in mind the conceptual distinctions among I-time, A-time, and B-time. To the extent to which researchers wish to appeal to B-times in connection with the question about causation or the question about vetoing, they should take steps to diminish the sources of inaccuracy.

ACKNOWLEDGMENTS

This article was completed during my tenure of a 2007-08 NEH Fellowship. (Any views, findings, conclusions, or recommendations expressed in this article do not necessarily reflect those of the National Endowment for the Humanities.)

NOTES