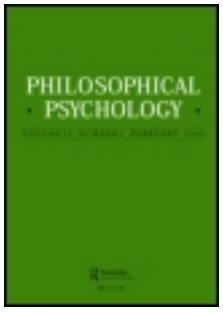

Philosophical Psychology
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/cphp20

Conscious intending as self-programming

Marc Slors [a]

[a] Radboud University, Philosophy , Erasmusplein 1, Postbus 9103, 6500 , HD , Nijmegen , The Netherlands

Published online: 31 May 2013.

To cite this article: Marc Slors (2013): Conscious intending as self-programming, Philosophical Psychology, DOI: 10.1080/09515089.2013.803922

To link to this article: http://dx.doi.org/10.1080/09515089.2013.803922

Philosophical Psychology, 2013
http://dx.doi.org/10.1080/09515089.2013.803922

Conscious intending as self-programming

Marc Slors

Despite the fact that there is considerable evidence against the causal efficacy of proximal (short-term) conscious intentions, many studies confirm our commonsensical belief in the efficacy of more distal (longer-term) conscious intentions. In this paper, I address two questions: (i) What, if any, is the difference between the role of consciousness in effective and in non-effective conscious intentions? (ii) How do effective conscious distal intentions interact with unconscious processes in producing actions, and how do non-effective proximal intentions fit into this process? I argue that answers to these questions point to a picture of distal conscious intending as a form of self-programming. The metaphor of "self-programming" will be elucidated by using a distinction between "structuring" and "triggering" causes. Though the self-programming metaphor does not amount to a full theory of conscious intending, I argue that it may be a useful heuristic in developing such a theory. I also argue that the metaphor is phenomenologically plausible.

Keywords: Causal Efficacy of Consciousness; Conscious Will; Intention

1. Introduction

In the past 30 years, psychologists and neuroscientists have produced a mass of evidence to the effect that conscious intentions do not causally initiate their intended actions. The issue here is not that consciousness cannot be an uncaused cause. No philosopher or scientist seriously doubts that every conscious activity in the brain is produced by preceding unconscious processes. But conscious activity might, in a colloquial sense at least, be said to initiate action if it coincides with the earliest reliable neural indicator of an upcoming action. Libet, Haggard, and others, however, have shown that this is not the case. They showed that simple conscious motor intentions occur only after the unconscious neural onset of actions (Brass & Haggard, 2008; Fried, Baaren, & Kreiman, 2011; Haggard & Eimer, 1999; Libet, Gleason, Wright, & Pearl, 1983; Soon, Brass, Heinze, & Haynes, 2008).

Marc Slors is Professor of Philosophy of Mind and Cognition at Radboud University Nijmegen.
Correspondence to: Marc Slors, Radboud University, Philosophy, Erasmusplein 1, Postbus 9103, 6500 HD Nijmegen, The Netherlands. Email: m.slors@ftr.ru.nl

Wegner (2002; Wegner & Wheatley, 1999) and others have argued that such intentions are not even part of the causal chains leading up to our actions. Moreover, the principles governing the unconscious roots of our actions are slowly coming into focus (Dijksterhuis, Bos, Nordgren, & Baaren, 2006; Dijksterhuis & Nordgren, 2006; Hassin, Uleman, & Bargh, 2004; Wilson, 2002). The picture that is drawn in this literature is that of conscious intending as an epiphenomenon; a causally irrelevant byproduct and/or rationalization of unconsciously produced actions. Although this picture is not uncontroversial (Baumeister & Masicampo, 2010; Bayne, 2006; Mele, 2006, 2009; Nahmias, 2010; Schurger, Sitt, & Dehaene, 2012), I will accept it, for the sake of argument, as a point of departure for this paper, but only insofar as it applies to so-called *proximal* conscious intentions, i.e., intentions that occur just before the action.

My main concern in this paper is with longer-term, or so-called *distal* intentions. These include e.g., my intention to take the nine o'clock train tomorrow morning, to stop before yellow traffic lights in the future, or to go to Italy on my next summer vacation. Although there certainly is a tendency to downplay the causal efficacy of this kind of intention as well in the literature, a large number of empirical studies, neatly summarized by Baumeister, Masicampo, and Vohs, show that "the evidence for conscious causation is profound, extensive, adaptive, multifaceted, and empirically strong" (2011, p. 351). Baumeister et al. do recognize, in agreement with Libet, Haggard, Wegner, and others, that "consciousness may be ill suited for direct control of physical behavior" (2011, p. 352). Thus, according to them, "conscious causation is often indirect and delayed, and it depends on interplay with unconscious processes" (2011, p. 331; see also Baumeister & Masicampo, 2010, p. 948). I have no quarrels with this position. But it does leave open two conceptual or philosophical questions, answers to which may contribute to charting the exact contribution of consciousness to our agency.

The first question is whether the difference between conscious intentions that are and those that are not effective is merely that—some (i.e., distal intentions) *are* effective, some (i.e., proximal intentions) not—or whether there is a further difference (besides the proximal/distal difference) that sheds light on the difference in causal efficacy. I will argue that there is a commonsensical distinction between two kinds of conscious intention that maps neatly onto the difference between effective and non-effective intentions. Thus, I will argue that the effective intentions of Baumeister et al. are a completely different class of mental process than the ineffective intentions of Libet, Haggard, and Wegner. Furthermore, I will claim that the non-efficacy of conscious short-term intentions is much more in line with everyday intuitions about ourselves than is usually suggested.

The second question is about the division of labor between effective conscious intentions and unconscious processes in producing behavior. How do effective conscious intentions "collaborate with" unconscious processes in producing behavior? Can we say more about the "dependency and interplay" of these two kinds of process, mentioned by Baumeister et al.? I will suggest that an important part of the answer to this question lies in recognizing the *kind* of causation that is involved in effective conscious intentions. Following Dretske (1988), I will distinguish triggering causes (roughly: the direct onset of a causal chain of events) from structuring causes

(roughly: the cause of the set-up that makes a given triggering cause have a specific effect). I will argue that it is crucial to see effective conscious intentions as *structuring* causes of our actions, while our actions are possibly always triggered unconsciously by external or internal stimuli.

The upshot of these answers is not a complete theory of all conscious and unconscious processes at play in the production of action. Rather, it is a picture of conscious intending that can best be captured in terms of a metaphor that may play a guiding role in the development of such a theory. This is the metaphor of self-programming: forming a conscious distal intention to do *x* in the future is attempting to program oneself to be unconsciously responsive to appropriate future stimuli in relevant contexts.

The paper is set-up as follows: I will start by briefly rehearsing the evidence for and against the efficacy of conscious intentions. I will then distinguish two radically different kinds of conscious intending that neatly map onto the effective/ineffective difference. Next, in order to account for the interplay between effective conscious intentions and unconscious processes in the production of actions, I will discuss the relevance of the distinction between triggering and structuring causes. Finally, I will evaluate the ensuing notion of conscious intending as self-programming and argue for its phenomenological plausibility.

2. Effective and Non-Effective Conscious Intentions

Psychological and neuroscientific attacks on the efficacy of conscious intentions take on different forms. In its weak form, the claim is simply that "the causal role of conscious thought has been vastly overstated" (Wilson, 2002, p. 107). This claim is relatively unspecific since the alleged causal role of conscious thought is not specified. A stronger and perhaps more common claim is that "behavior does not originate with a conscious decision" (Dijksterhuis, Chartrand, & Aarts, 2007, p. 52). The idea here is not so much that consciousness has nothing to do with the processes that causally lead up to our actions. Rather, the idea is that the initiation of such processes is unconscious. This claim harks back to Libet's experiments, which demonstrated that conscious intentions to perform simple finger-lifting movements occur after the unconscious neural onset of such actions (Libet et al., 1983). Such experiments have been replicated in various forms (Brass & Haggard, 2008; Fried et al., 2011; Haggard & Eimer, 1999; Soon et al., 2008).

It is important to note that these claims leave room for some role for consciousness in the production of action. It is one thing to claim that the onset of actions is unconscious, but it is another thing to claim that consciousness is completely irrelevant in the process leading up to an action. Libet himself was keenly aware of this. He showed, experimentally, that consciousness does have the ability to veto unconsciously initiated actions (Libet, 1985). On Libet's picture, although consciousness does not initiate the causal chain of neural events leading up to an action, it still is a part of that chain, and as such, it can exert some influence. A similar line of thought leads Alfred Mele to conclude

that consciousness may indeed be even more important than Libet makes it out to be. On Mele's picture, the unconscious initiation of an action should not be equated with an intention or decision to act, but merely with an initial urge to act. It is only when the process leading up to an action reaches its conscious stage that an intention is formed or a decision is taken to act, by not vetoing the urge and hence by consciously letting the action occur (Mele, 2006, 2009).

Such room for conscious intervention is eradicated in the strongest form that scientific attacks on the efficacy of conscious intentions can take. On this form, conscious episodes preceding actions are neither the initiators, nor even part of the causal chain leading up to an action. They are side effects of the (unconscious) processes that start such a chain. Since consciousness of the upcoming action, or of thoughts and/or intentions that pertain to the action, precede the actual physical movements, we are presented with the illusion that consciousness causes our behavior. In reality, it is claimed, there is merely temporal succession and no causation. The most outspoken defender of this model is Wegner (2002, 2005; Wegner & Wheatley, 1999). But, precisely since none of the major attackers of the causal efficacy of conscious intentions appears to be set on leaving a role for consciousness (with the exception of Libet), this model may well be what implicitly underlies most rejections of the efficacy of conscious intentions. For the sake of argument, therefore, I shall take Wegner's epiphenomenalism as the exemplary form that the scientific rejection of the efficacy of conscious intentions can take.

I will not question Wegner's model as far as proximal intentions are concerned. But does it apply to distal or longer-term intentions? Wegner suggests, at times, that it does. However, Baumeister et al. have summarized a wealth of evidence to the effect that consciousness does have causal influence on our actions in the long run (2011; Baumeister & Masicampo, 2010). Much of this evidence does not specifically pertain to conscious longer-term intentions, though. Thus, e.g., a lot of attention is devoted to the fact that mentally simulating actions increases the likelihood of those actions being carried out in the future. Likewise, replaying, interpreting, and reflecting on past events clearly influences subsequent actions. But such causal influence does not show unequivocally that conscious intention formation is causally involved in the production of actions later on. Evidence to that effect *is* available, however.

Most prominent among this evidence is the work on so-called "implementation intentions." Dozens of studies confirm the notion that intentions are very significantly more likely to be carried out when agents consciously intend to carry out specific actions rather than to pursue more abstractly formulated goals (Gollwitzer, 1999). Thus, for example, when women who already held the goal of performing breast self-examinations were asked to consciously form implementation intentions, 100 percent of them reported to have actually done so, while only half of the women who were not asked to form implementation intentions reported having performed self-examinations. Various similar studies have yielded similar results (Gollwitzer & Sheeran, 2006). Indeed, the conscious formation of implementation intentions is in the long run a better guarantee for actually carrying out the intended actions than cue-behavior association procedures (Papies, Aarts, & De Vries, 2009).

In other studies, conscious longer-term intention formation is induced by letting subjects anticipate regret. Such anticipation yields vigilance and careful information gathering, and also promotes risk avoidance (Janis & Mann, 1977; Richard, Van Der Pligt, & De Vries, 1996; Tetlock & Boettger, 1994). It can make people choose the more defendable options (Simonson, 1992). While in some cases, anticipating regret can induce people to avoid making decisions (Anderson, 2003; Beattie, Baron, Hershey, & Spranca, 1994), risk avoidance, vigilance, and opting for defensible choices clearly shows the influence of intentions formed due to consciously anticipated regret on future actions.

While empirical evidence for the efficacy of conscious distal intentions is important, I believe that simple everyday-life experience is at least an equally important factor in resisting Wegner's model with respect to longer-term conscious intentions. When I form the intention, while checking the airline website and making the relevant reservation, to take the 2:15 flight to London next Wednesday, I usually find myself in that airplane on that day. I am well aware of the fact that an appeal to everyday experience is tricky, specifically in view of the fact that Wegner's model does acknowledge the *experience* of the causal efficacy of conscious intentions. But there is a crucial difference between the kind of conscious intending involved in the experiments Wegner bases his model on and the kinds of distal intending that everyday experiences such as the one described above suggest to be efficacious. I will outline this difference in the next section and explain why distal conscious intending *cannot* fit Wegner's epiphenomenalism.

3. Two Kinds of Conscious Intention: The Limits of Wegner's Model

There is a striking tendency in a significant portion of the neuroscientific and psychological literature on conscious intending (as well as in the philosophical commentaries on that literature) to consider the class of conscious intentions to be more or less homogeneous. To be sure, proximal intentions are distinguished from distal ones, and simple motor intentions are distinguished from complex, more abstract ones. But the fact that such various forms of intending can be conscious is very often taken to denote one specific feature that intending processes, of whatever kind, might exhibit. This, I contend, is a mistake. There are at least two different ways in which intending can be conscious. 'Consciously intending' can mean (1) becoming or being conscious of the fact that one intends to do *x*, or it can mean (2) consciously forming the intention to do *x*. The difference hinges on the passive/receptive versus the active/formative role of consciousness.

The functional difference between consciousness as passively registering and as actively producing intentions is central to my argument. But given that this difference coincides more or less perfectly with a widely acknowledged difference in kinds of consciousness, I will elaborate on this difference by introducing these kinds.[1] There are many ways in which the class of conscious mental states can be subdivided into sub-classes, but I shall stick to a widely accepted two-part subdivision. This subdivision is

derived from the distinction between perceptual and reflective consciousness (Roessler & Eilan, 2003). Perceptual consciousness is the conscious registering of states of affairs in the world or in one's body, usually via the senses. Reflective consciousness is the ability to reason and to reflect on one's experiences. That distinction is intuitive, but not entirely suitable for the case of intentions, for the simple reason that one's own intentions are not objects of sensory perception. We can, however, broaden the category of perceptual consciousness to the category of phenomenal consciousness, which includes perceptual consciousness as well as other subjective experiences, such as moods or dreams (Block, 1995, 2005). Phenomenal consciousness is the ability to become subjectively aware of feelings, sensations, and perceptions. What we end up with is a distinction between reflective and phenomenal consciousness. Much more can be said about this difference, but I believe the difference between the two kinds of consciousness is intuitive enough to proceed.

With respect to the different ways in which intentions can be or become conscious, the point is that conscious reflection can actively produce intentions, while phenomenal awareness of intentions merely consists in the passive registration of the fact that one has this or that intention. Let me go over both kinds of intending in a little more detail, starting with the latter.

Being or becoming conscious of an intention one has requires phenomenal awareness, but not reflection—that is, it requires registering rather than formation. Libet's finger-lifting intentions are excellent cases in point. There is no way, really, one can reason or reflect about when to lift a finger in these experiments. For there are simply no *reasons* to lift a finger now rather than two seconds later. Thus, we cannot speak of a decision to lift one's finger in the most common sense of that word, i.e., the sense in which one decides between options based on considerations. The "decision" to lift one's finger now is not a product of conscious reflection; one simply becomes conscious *of* the fact that one wants to lift one's finger *now*. Or rather, that one is about to lift one's finger. The intention to move is consciously registered before the movement is executed. But, and this is crucial, it is not *issued by* consciousness.[2] It is registered consciously—one is phenomenally aware of the intention—but the intention itself may very well have been issued or produced unconsciously, which is exactly what Libet found.

My contention that consciousness is passive with respect to the finger movement intentions involved in Libet-style experiments conflicts with Mele's (2006, 2009) reading of what happens in such experiments. Mele stresses that although the urge to move one's finger is indeed unconsciously produced, conscious affirmation of this urge is required to actually produce a finger movement. Consciousness turns a mere urge into a real intention. In other words, consciousness is not passive on Mele's picture. However, Mele's defense of the efficacy of proximal conscious intentions assumes that the conscious episode just prior to the finger movement is part of the causal chain that leads from an unconscious urge to an action. This is the model Libet sketches. But this model is not compatible with Wegner's epiphenomenalism, which views conscious intentions not as part of an ongoing causal chain, but as dead ends that branch off of the causal chain leading from unconscious causes to actions. I have

taken Wegner's epiphenomenalism as my point of departure, for the sake of argument, since it presents a worst-case scenario—consciousness is not causally connected to the finger movement and can therefore simply not fulfill the role Mele envisages.

The idea that we sometimes become passively conscious of intentions we happen to have does not strike me as odd or out of line with common experience. I sometimes become conscious of my thirst and my intention to quench it when I am already on my way to the fridge. Other examples include intentions to make certain tactical movements in a sports match, or, if you are a musician, intentions to play certain notes or riffs in an improvisation.[3] In such cases, there simply is no time for conscious reflection.

Intentions that we become conscious of, but that we don't consciously produce, are not only involved in Libet's experiments. They are typically also involved in Wegner's own most important study, the "I Spy experiment" (Wegner & Wheatley, 1999). In this experiment (which is too complex to discuss in detail here), thoughts are "inserted" into the heads of people just prior to their being coerced, unbeknownst to them, to perform an action that matches the thought. In such cases, people turn out to be inclined to think that their action was caused by their thought. Wegner concludes from this that a conscious thought matching an action and preceding that action is perceived as being the cause of that action, even when it demonstrably is not (the cause of the action is cleverly disguised coercion). What is important is that the conscious thought at play here is one that subjects become conscious of—passively—and not one that they themselves actively and consciously produce.

In fact, intending through conscious reflection is virtually absent in Wegner's discussion of what he calls his "theory of apparent mental causation." It surfaces only in his discussion of confabulation, that is, of experiments that show how we are inclined to fabricate reasons for our actions after the fact (Gazzaniga, 1988, 1995; Nisbett & Wilson, 1977). Such reason giving often takes the form or recounting the fictitious conscious reflection prior to an intention. Interesting as these cases may be, here I am concerned with real intentional conscious reflection. While absent from Wegner's discussion, this is typically the kind of conscious intending that figures in the studies that show the efficacy of distal conscious intending. Gollwitzer's implementation intentions are not passively consciously registered; they are actively consciously produced. Likewise, intentions that are formed as the result of anticipating regret are the product of conscious reflection.

This is what sets the class of intentions that are produced by conscious reflection apart: they do not exist prior to conscious reflection and would not have existed without conscious reflection. This is not to deny that all conscious intention formation is initiated by unconscious processes. The point is, rather, that the process of conscious reflection, however unconsciously initiated, results in an intention that in many cases could not have already been there in unconscious form. When Gollwitzer asks his subjects to form a certain implementation intention, then it is in response to his request that subjects form such an intention. Before hearing and considering this request, there cannot have been an unconscious version of the implementation intention. As another example, consider my consciously planning to take the 2:15

flight to London on Wednesday. Supposing that I plan a flight like this for the first time, so that I cannot rely unconsciously on previous experience or unconscious routines, I can only form that intention when I have consciously checked the airline website. Before my process of conscious intention formation, I could simply not have had the intention to take the 2:15 flight; I did not know of the existence of such a flight before I checked the site—consciously.

If consciously formed intentions cannot exist prior to the event of intention formation, Wegner's epiphenomenalism cannot apply to them. According to Wegner's model of apparent mental causation, the unconscious processes leading up to the conscious intention and the unconscious processes leading up to the consequent action are identical, or at least connected. Hence, there is complex unconscious activity that first produces a conscious intention and then leads to an action. The action matches the intention but is not the effect of it; it is the effect of the common unconscious cause. If this common unconscious cause is indeed the origin of the action—that is, if there is no influence from the conscious intention on the production of the action—and if the action does indeed match the intention, then the content of the intention must already be "present" in this common unconscious cause. This may indeed well be what happens in the case of intentions of which one becomes conscious. But it cannot be what happens with intentions that are consciously formed. For before the actual process of conscious reflection, it is uncertain what the ensuing action must look like, as the examples above show.

Since much—i.e., the non-applicability of Wegner's epiphenomenalism to consciously formed distal intentions—hinges on the correctness of the idea that some intentions are indeed formed during and not prior to conscious reflection, let me defend that idea in more detail against possible skeptical doubts. For it might be argued that the distinction between becoming (phenomenally) conscious of an intention and (reflectively) consciously forming an intention is part of our commonsense psychology rather than a real scientific distinction. Moreover, it seems not unlikely that unconscious processes may be responsible for the kind of elaborate reasoning that we normally assign to reflective consciousness. Given these considerations, could it not be the case that what I have described as conscious intention formation is simply becoming conscious *of* the process and outcome of an unconscious reasoning routine? If so, Wegner's model may still apply, despite appearances to the contrary.

This objection may even apply to the example of consciously deciding which plane to catch by checking the airline website. This example is chosen precisely because it seems to be the case, at first glance at least, that no conscious decision can be taken before the information from the website is consciously processed. If that is correct, this would block the notion that my conscious decision could have been prepared at an unconscious level. What the objection alluded to above boils down to, however, is that when we zoom in on the details of this process, there may be room for a Wegner-like depiction of what goes on after all. For, (1) it may be the case that the information from the website is in fact unconsciously processed and (2) fed into an unconscious reasoning mechanism which arrives at a decision *of* which I become conscious (while it appears to me as if I consciously *form* the decision).

Let me admit straightaway that there is no *logical* argument against this option. There are, however, good scientific reasons to think that this is not what is going on. The most powerful, in my opinion, stems from the near consensus among consciousness researchers over the function of (both kinds) of consciousness. Referring to eminent researchers such as Baars, Crick, Edelman, and Deheane, Baumeister and Masicampo stress that "phenomenal awareness and conscious thought, enable the different parts of the mind and brain to share information with each other .... Consciousness has been seen by almost all theorists as helping to integrate information" (Baumeister & Masicampo, 2010, p. 949). Morsella refers to a much longer list of theories on the function of consciousness when he speaks of "the integration consensus." The idea behind this consensus is that conscious processes "integrate neural activities and information processing structures that would otherwise be independent .... Many of these theories speak of a central information exchange, where dominant information is distributed globally" (Morsella, 2005, pp. 1001–1002).[4]

Applied to the example of checking the airline website in order to decide which flight to catch, the idea here would be that the function of *consciously* processing the information from the website is precisely to "broadcast this information globally" across the brain so that it is available to various subsystems, including the system that will make the decision based on this information. Consciously processing the available options is a means of "laying all the options on the table," allowing decision-making mechanisms to access all the relevant information. Assuming that the website information may well be processed unconsciously is failing to provide an explanation of why and how this information is combined with one's preferences and other thoughts and how the total package of information reaches the relevant decision-making subsystems of the brain. Assuming that the information is consciously processed does provide a widely accepted explanation that, moreover, fits extremely well with our experience.

It is reasonable, then, to assume that consciousness *is* needed in the first stage of the decision-making process—it is the means of getting all the relevant information to the right places. There is much more to be said about the function of consciousness after this initial stage. Thus, the decision-making or intention formation on the basis of considering available options might well involve imagining possible future scenarios. This requires consciousness.[5] But focusing on the first stage is enough to make my point: the decision of which flight to take cannot have been formed prior to the decision-making process that involves consciously laying the options on the table, so to speak. And that means that Wegner's model cannot apply to this kind of conscious decision-making or intention formation.

The "trick" of the website example is, of course, that outside information has to be processed (consciously) before an intention can be formed. This rules out the possibility that the intention has already been prepared at an unconscious level. But many consciously formed distal intentions do not rely on processing outside information. Suppose I am in doubt about whether our next vacation should be spent in Spain or in Italy. At a given point in time, I sit down and weigh all the pros and cons against each other and make a decision. This could be a consciously formed decision.

Of course, we should not rule out the possibility that my decision was already made at an unconscious level because I had an unconscious preference for Italy. In that case, my weighing pros and cons was nothing but a rationalization. But neither should we rule out the possibility that the decision was actually made in a conscious process of weighing the options. And, like in the website example, the strongest argument for the necessary involvement of consciousness in that process is its "global broadcasting" function. By becoming conscious of the pros and cons, all options are being put on the table; they become accessible for decision-making subsystems of the brain. Whether or not the actual decision-making is conscious, it is at least the case that a necessary part of the process of intention formation (laying the options on the table) must be conscious for all the relevant information to be brought together and be accessible to the relevant parts of the brain. If my consciously deciding to go to Italy was not a rationalization of an unconscious preference but a real decision, then it seems plausible that consciously laying the options on the table was a conditio sine qua non. If so, it could not have been prepared unconsciously, and Wegner's model does not apply.

4. Triggering versus Structuring Causes: The Idea of Self-Programming

Noting that conscious distal intention formation is, at least in many cases, not a causal epiphenomenon is not the end of the story. In fact, it is only the beginning. For the causal efficacy of conscious distal intention formation raises important questions.

In order to see what questions are raised, it is good to start by recognizing that there can be no simple, straight causal arrow from a process of distal intention formation to an action or series of actions. Pacherie (2006), among others, describes the process that leads from a distal intention to an action or series of actions as a cascade of intention-transitions. Distal, future-directed intentions are turned into proximal, present-directed intentions in a process that Pacherie calls "situational anchoring." For instance, my intention to take the 2:15 flight to London on Wednesday must, in order to be carried out on that day, lead to an action plan that gives rise to specific present-directed, proximal intentions, e.g., the intention to print out my boarding card, to take a cab to the airport, to get through customs, etc. These proximal intentions, in turn, must give rise to specific motor intentions that, in the end, cause the relevant actions. The details of such a dynamic view on intentions are intricate, and many of them are still debated. Here I will only focus on the intentions that occur just before one performs the actions that constitute carrying out one's distal intention—be they proximal or motor intentions. I will refer to them as "anchoring intentions."

Pacherie is not explicit about the conscious or unconscious nature of anchoring intentions. Nor is she explicit about the mechanism that initiates the anchoring process. One view on these issues that may easily be read into her dynamical theory of intentions is this: following Brand (1984) and Mele (1992), Pacherie stresses that one of the functions of intentions is to monitor the progress toward achieving the intended goal. Hence, as a consequence or part of the formation of a distal intention, a person may be thought to monitor (presumably not consciously all the time) whether the

appropriate situation arises for carrying out the distal intention. When appropriate circumstances are detected, a person becomes conscious of this and translates the distal intention into a (series of) conscious anchoring intention(s). These intentions, in turn, cause our actions.

It is not important whether this view on the mechanisms behind intentional anchoring is actually held and defended by anyone explicitly. But it is instructive to see how this position should be replaced with a different picture in view of the discussion of the previous sections. For, however intuitive the picture sketched above may seem to some, it is not compatible with the point of departure of this paper: the presumed correctness of Wegner's epiphenomenalism with respect to conscious proximal intentions. On Wegner's picture, the conscious proximal intentions that may result from the anchoring process of a distal intention are only the apparent causes of our actions. The real causes would have to be unconscious anchoring intentions. We may become *conscious of* these anchoring intentions (see the previous section), but that consciousness is in fact ineffective (for important exceptions, see the next section).

If this sounds absurd, consider the following small-scale version of a similar phenomenon. Subjects who are asked to press a button as soon as they feel a tap on their arm usually respond in about 100 milliseconds. But they become conscious of responding only after 500 milliseconds (Wegner, 2002, p. 56). The idea here is that a similar phenomenon is at play in the situational anchoring of distal conscious intentions: we consciously form our distal intentions but consciousness only plays a registering role—not a formative or active one—in the processes that transform distal intentions into anchoring intentions. The idea that situational anchoring is not facilitated by consciousness is in fact underscored by experiments that show that consciousness need not even be involved. Bargh, Lee-Chai, Barndollar, Gollwitzer, and Trötschel (2001) have devised a series of experiments in which consciously formed goals were unconsciously triggered and run to completion, attaining the desired outcomes.

The non-active involvement of consciousness in intentional anchoring is illustrated by William James through the following example:

> We all know what it is to get out of bed on a freezing morning in a room without a fire, and how the very vital principle within us protests against the ordeal. Probably most persons have lain at certain mornings for an hour at a time unable to brace themselves to the resolve. We think of how late we shall be, how the duties of the day will suffer; we say "I *must* get up, this is ignominious," etc.; but still the warm couch feels too delicious, the cold outside too cruel, and resolution faints away and postpones itself again and again .... Now how do we *ever* get up under such circumstances? If I may generalize from my own experience, we more often than not get up without any struggle or decision at all. We suddenly find that we *have* got up. (2001, p. 291)

Crucially, James thought that this simple example contained "the data for an entire psychology of volition" (Richardson, 2007, p. 280).

Once the causal efficacy of conscious anchoring intentions is deleted from the intuitive (straw man) picture of the intentional cascade, there is reason to make a further adjustment. If intentional anchoring is an unconscious process, the concept of

a monitoring process issued by distal intentions that actively and consciously "decides" whether circumstances are appropriate to turn a distal intention into matching anchoring intentions becomes less attractive as well. For the idea that consciousness is efficacious in intentional anchoring is connected with the idea that one recognizes, consciously, that the appropriate circumstances have arrived to translate distal intentions into anchoring intentions. If intentional anchoring is unconscious, by contrast, it may well be that anchoring intentions are formed in *direct* response to specific stimuli.

Take my consciously formed intention to stop at yellow traffic lights (rather than step on the gas, like I used to do). The anchoring of this intention is triggered by the stimulus of perceiving a yellow traffic light. None of this triggering needs to involve much conscious thought, and there is nothing in my phenomenology that even remotely resembles my monitoring of situations after I had made my conscious decision to stop at yellow traffic lights from now on. Here it could be asked, however, why it may not be the case that my conscious decision issues an *unconscious* monitoring process that triggers the anchoring processes when the relevant circumstances are detected. The best answer here, in my view, is simply to insist that there is a "leaner" option.

Distal intentions must surely cause the fact that we are responsive to very specific triggers in very specific situations. If this responsiveness is not the result of a monitoring mechanism that selects relevant triggers in relevant circumstances, then how do distal intentions contribute to the process of intentional anchoring? Here the answer is to replace the single kind of causation at play in the intuitive (straw man) picture of Pacherie's intentional cascade with two *different* kinds of causation. On the intuitive (straw man) picture, conscious distal intentions cause a (partly conscious) monitoring mechanism to be active, which, when the appropriate circumstances arrive, cause conscious anchoring intentions, which in turn cause actions. An alternative story that fits well with the idea of unconscious anchoring would be to say that conscious distal intention formation causes me to be disposed to be responsive (in the right kind of way) to those triggers that set off the intentional anchoring process—be they external as in the traffic light example, or internal (and hence probably intractable) as in James' example or Libet's finger-lifting. No active monitoring is needed.

A distinction between two kinds of causation, introduced by Dretske (1988, p. 43), may be of help here. Dretske distinguishes triggering causes from what he labels "structuring causes." Normally, when x causes y, x is the triggering cause of y (there is an extensive philosophical literature on what *that* is supposed to mean, but there is no need to go into that here). But there may also have been a process in which z causes the fact that the occurrence of x will produce y. Here, z is a structuring cause. A concrete example might help. My flicking the switch in my study is the triggering cause of the fact that the light in the study goes on. But it can only be the triggering cause of the light going on given the fact that the switch and the lamp are appropriately connected (to each other and to the electricity net). The actions of the electrician who originally wired them up appropriately are the structuring cause of the event of my flicking the switch causing the light in my study to go on.

Conscious distal intention formation is like the electrician's activity; it is a structuring cause. It causes the fact that certain internal or external stimuli will cause specific actions. Using a metaphor, we might say that it causes us to be "programmed" to be responsive in specific ways to specific stimuli in specific circumstances. Conscious intention formation is a form of self-programming.

The self-programming metaphor paints a different picture of dynamic intentions than the intuitive (straw man) version of Pacherie's intentional cascade. On this alternative picture, there is no series of causal arrows from distal intention to monitoring mechanism, from monitoring mechanism to anchoring intention and from anchoring intention to action. In fact, on this alternative picture, talk of intentions as causally efficacious items in our heads is metaphorical.[6] We may describe a reflectively conscious process as the conscious formation of an intention, arriving at a conclusion or making a decision. But such an intention, conclusion, or decision *is* in fact a complex change in one's behavioral dispositions.[7] Once this dispositional change has occurred, anchoring is a process of responding unconsciously to internal or external triggers as a consequence of one's changed dispositions. It is not a process in which intentions (as causally efficacious items) are transformed.

This is what Baumeister et al. hint at but do not spell out when they say that "conscious causation is often indirect and delayed, and it depends on interplay with unconscious processes" (2011, p. 331). The notion of self-programming explains why conscious causation is indirect: processes of "conscious intention formation" are structuring causes, but not the triggering causes of actions. It also explains why conscious causation is delayed: after our self-programming, the actual triggering of the action still has to happen. Finally, the interplay between conscious causation and unconscious processes Baumeister et al. mention is neatly captured: conscious intention formation is initiated unconsciously, but it yields new intentions (i.e., intentions that were not there in unconscious form before). These intentions program a person to be responsive to specific stimuli in specific circumstances; when these circumstances occur, the situational anchoring of the intention is unconsciously triggered by the appropriate stimuli. This may lead to the person becoming (passively) conscious *of* the actual present-directed intentions or motor intentions behind the actions that carry out the consciously formed intention.

The self-programming metaphor fits extremely well with the finding of Gollwitzer (1999; Gollwitzer & Sheeran, 2006) and others that so-called implementation intentions give a significantly higher chance of reaching one's goals than mere goal intentions. Whereas goal intentions only specify the goal that one intends to reach (e.g., losing weight), implementation intentions specify the behavior one will perform in the service of achieving that goal and, crucially, the situational context in which one will perform the relevant actions (e.g., in the case of losing weight, taking only one sandwich from the cafeteria at lunchtime). Implementation intentions have an if-then format. By specifying the behavior needed to achieve one's goals and the situations in which it needs to be performed, formation of implementation intentions constitutes an *explicit*, and hence stronger, form of self-programming than the formation of mere goal intentions. The idea of implementation intentions seems to arise from the tacit

recognition that distal intention formation is a form of self-programming that can in fact be made more effective by becoming more explicit.

5. Two Caveats

I would like to be clear about the way in which, and the extent to which, I think that the notion of conscious intending as self-programming is explanatorily useful. Two caveats are in order.

First, the notion is a general *model* that requires further differentiation and detailing in concrete cases. I have deliberately simplified the discussion for the sake of clarity of exposition by ignoring the possibility that the dynamic development of consciously formed distal intentions may at some point involve the further *active* conscious formation of plans of action that are in the service of carrying out the overall intention. An example may help here. If I carry out my intention to be on the 2:15 flight to London by going through customs an hour earlier, I may need to form some conscious intentions in order to do so. I may need to remind myself, consciously, of the fact that I must not have liquids with me in containers larger than 50 ml, and that I need to have them bagged (such conscious thoughts may also be a response to a sign prompting me to check my liquids). This may lead to forming the intention to check my liquids, get the required plastic bags, and make sure my toothpaste and shampoo are appropriately packaged. Even if I act almost immediately upon thinking I should check my liquids, this is still a small-scale form of self-programming on the model proposed. For the exact actions I perform in order to carry out my intention are determined by the factual circumstances I am in (if I am in a line, I might wait until I am in a position to put my suitcase down, e.g., on a table).

Two things are important to notice. First, by allowing for intentions that are relatively proximal to actions to be small-scale structuring causes of our actions, the model proposed can allow for many intentions that are intuitively classified as proximal intentions to be causally efficacious in producing actions. The distinction between efficacious and non-efficacious conscious intention formation is best captured in terms of the difference between structuring and triggering causes, and not in terms of the relatively imprecise distinction between distal and proximal intentions I started out with. Secondly, in concrete cases, the model of a structuring cause followed by an unconscious anchoring process in which actions are triggered by internal or external stimuli allows for similar, "nested" processes, which considerably complicates the interplay between conscious and unconscious processes.

The notion of conscious intending as self-programming, both in simple and complex cases, is not intended as a concrete theory of conscious agency. It is an abstract idea. But its explanatory value is precisely in the fact that it can be applied at the level of cognitive architecture. As such, it can be a useful heuristic for further cognitive neuroscientific research into the mechanisms behind the conscious causation of actions. Thus, I claim that the idea of conscious intending as self-programming contributes to the solution of the problem of "mental causation," as

Wegner often calls it. Wegner ignores the possibility of structuring causes, and that is exactly why he cannot see "mental causation" to be anything but an illusion.

But here, a second caveat is in order. For the notion of "mental causation" is used in subtly different senses in different areas of scientific and philosophical discourse. What the notion of self-programming explains is how certain brain processes that "underlie," "constitute," or "realize" conscious processes of intention formation can play a causal role in the production of actions. What the notion does *not* explain is how such processes contribute to the causal processes leading to actions *in virtue of* their mental or conscious properties, *rather than* their neurobiological properties. This is the domain of *another* debate on mental causation, one that is exclusively within the domain of philosophy (Heil & Mele, 1993; Kim, 1998, 2005). This debate is orthogonal to the issues discussed in this paper, and I do not claim to contribute to it in this paper.

6. Self-Programming As a Realistic Model of Conscious Intentional Agency

The idea that consciously forming distal intentions is a form of self-programming explains not only why Baumeister's examples of the causal relevance of conscious intentions are perfectly compatible with Libet's and Wegner's claims about the inefficacy of proximal conscious intentions, it also explains how the two are related. The idea that distal intention formation is not a triggering but rather a structuring cause behind our actions explains why and in what sense conscious causation of action is delayed, indirect and dependent upon interplay with unconscious processes. But even in view of the fact that Wegner overlooks the notion of structuring causes, which is an omission he may admit, he may still think that the idea of conscious intending as self-programming is not enough "to give us what we want," that is, to give up the idea that conscious causation is an illusion. For Wegner argues that the phenomenology of conscious intentional agency matches what James has labeled the "ideomotor theory." This is the idea that conscious intentions preceding actions are—under specific conditions (Wegner, 2002, pp. 63–98)—experienced as the triggering causes of actions. This experience would still be contradicted by the self-programming model I advocate. For on this model, distal intentions are "merely" structuring causes and proximal intentions are not triggering causes of actions.

In this final section, however, I will argue that the ideomotor theory is misguided as far as the phenomenology of proximal intentions is concerned. I will also argue that the idea that distal intentions are "merely" structuring causes of our actions fits very well with the phenomenology of forming and executing conscious distal intentions. Hence, the self-programming model provides a phenomenologically realistic picture of conscious intentional agency. It gives us "everything we want," i.e., every reason to reject the idea that the science of conscious intentional agency shows us that we live in an illusion.

As to proximal intentions, my first argument against the ideomotor theory (as far as *conscious* intentions are concerned) relies on the distinction between the two types of conscious intentions discussed in section 2. Few proximal intentions are truly

consciously *formed* intentions (as opposed to intentions we become conscious of). By consciously formed intentions, I mean intentions that would not have been there without conscious reflection or reasoning. The reflection or reasoning has to be such that it produces, rather than registers or notices our intention. With proximal intentions, this is rarely the case. Thus, for example, when I am consciously deciding whether I will buy vanilla or strawberry ice cream and finally opt for vanilla, this does not count as a consciously formed intention. For there are no *reasons* to choose vanilla other than my becoming aware *of* my preference. It may well be the case, and very probably is the case, that some unconscious urge to opt for vanilla was already present before my conscious "deliberation." I just had to become aware of it. We shouldn't even rule out the possibility that scientists may have been able to predict my choice before I was conscious of it.

The point is that this does not contradict our everyday experience. It seems perfectly commonsensical that in consciously "deciding" between vanilla and strawberry ice cream, I am doing nothing other than figuring out an already present but not yet conscious preference. Thus, the idea that my previously unconscious preference is ultimately pulling the strings does not uncover an illusion; it doesn't contradict our everyday thinking about ourselves. Likewise with Libet's finger-lifting. I take it that most subjects would agree that at some moment, they just felt an urge to move, of which they became conscious prior to moving. The notion that we experience such conscious proximal intentions as causes of our actions stems largely from Wegner's (Wegner & Wheatley, 1999) discussion of his I Spy experiment in which subjects judge actions to be caused by them when their prior thoughts match the ensuing action (even though there is no causal connection). However, what these subjects reported is that they considered these actions to be *their own doing*. And an action can just as much be experienced as my own doing when I know it to be issued by *my own* unconscious urges of which I became conscious just prior to acting.[8] Wegner, like many other scientists and philosophers, seems to have a blind spot for the fact that in ordinary life we consider unconscious preferences and urges to be our own just as much as our conscious thoughts.

My second argument against the ideomotor theory of conscious proximal intentions relies on research into the phenomenology of intentional agency. Briefly put, many studies show that the ideomotor model (as far as conscious proximal intentions are concerned) is an incorrect, idealized description of our experience of intentional action. In fact, James already showed this in the example cited above. But present-day studies show that many actions we consider to be intentional are simply not preceded by conscious intentions in the way the ideomotor model would predict (Bayne, 2006, 2008; Horgan & Tienson, 2002; Horgan, Tienson, & Graham, 2003). Thus, all in all, the fact that the self-programming model is incompatible with the ideomotor model of conscious proximal intentions does not show that it contradicts our phenomenology.

As to distal intentions, the notion of structuring causes and self-programming provides an elegant account of a salient feature of the phenomenology of conscious distal intentions. Not all consciously formed distal intentions are ultimately carried

out. Sometimes we forget our intentions; other times, we are simply weak-willed. My claim is that the phenomenology of weakness of will is particularly well captured by the self-programming model. On this model, weakness of will is a failed attempt to program one's own largely unconscious motivational system. The motivational system, in other words, has a "will of its own." And that, I submit, is exactly how we experience, from the inside, attempts to be strong willed: we have to counter "forces" in ourselves that we do not want to identify ourselves with (Frankfurt, 1971, 1988). But not by finding ourselves in a constant conscious battle. When we find ourselves to be weak-willed, we have to conclude after the act (often time and time again) that we did not act as intended. The "battle" itself is fought mostly at the unconscious level.

On a model of the dynamic development of consciously formed distal intentions in terms of the monitoring of the progress made in carrying out intentions, e.g., the intuitive (straw man) model of section 4, weakness of will must at some point involve an active conflict between processes that initiate the anchoring of the distal intention due to registering the appropriate circumstances and processes that resist this anchoring. The point here is that we rarely experience such a conflict in cases of weakness of will. Of course, it may be argued that the conflict is at an unconscious level. This would be allowing for the phenomenology of weakness of will by screening off our experience from the processes that drive our behavior.

On the self-programming model, however, there is no need to invoke an active conflict in the processes that drive our behavior in order to explain weakness of will. On the self-programming model, weakness of will is failing to make oneself responsive to relevant stimuli in appropriate circumstances. This, I submit, is how we experience being weak-willed. And in this respect, the idea that forming a distal conscious intention is attempting to provide the structuring causes for certain actions is not just compatible with a crucial aspect of the phenomenology of the human condition (which can in fact be said of most models when screening off phenomenology from subpersonal behavior-driving mechanisms is an option), but it actually captures this aspect.

All in all, I conclude that the idea that consciously formed distal intentions are not triggering causes of actions, but "merely" structuring ones, fits the phenomenology of conscious agency remarkably well.

Acknowledgements

Thanks to three anonymous referees for this journal for extremely helpful comments on earlier drafts of this paper.

Notes

[1] Baumeister does recognize the different kinds of consciousness involved (Baumeister & Masicampo, 2010, p. 945; Baumeister, Masicampo, & Vohs, 2011, p. 3). However, he does not highlight the fact that this difference is responsible for the difference in kinds of conscious intentions outlined in this section.

[2] The premise here is that for my intention to be issued or produced consciously, I must be able to tell how this process developed. After all, I was consciously aware of it. It seems reasonable to require some sort of reason-giving here, however minimal. Merely stating that at time *t* I suddenly felt like moving my finger should count as reporting the passive conscious registration of an intention at time *t*.

[3] There is more to be said about the musical improvisation example. Baumeister, DeWall, and Vohs (2008) have shown that although improvisation may start unconsciously, consciousness does seem to be required for truly creative improvisation. The example is taken from my own experience, and the findings of Baumeister et al. are compatible with that. In my experience, consciousness during improvisation often plays only a monitoring role and starts playing a more invasive role when things go wrong or when one's notes become dull. Thus, in reality, improvisation involves a complex interplay between unconscious and conscious processes. The example here pertains only to those parts of an improvisation where musicians "hear themselves play."

[4] Morsella speaks explicitly of phenomenal consciousness only. His definition of phenomenal consciousness, however, differs from mine above precisely in that phenomenal consciousness, on his understanding, is not restricted to passive registering. From the literature he cites when describing the integration consensus, it is clear that what I have labeled reflective consciousness also falls under this consensus.

[5] For an overview of the working and function of conscious simulation and other relevant modes of conscious thought and their functions, see Baumeister and Masicampo (2010).

[6] In fact, I agree with Ryle (1949) that thinking of intentional action as action caused by intentions is a category mistake that is derived from a "paramechanical" Cartesian view of the mind-body relation.

[7] The complexity of the change is due to the fact that most distal intentions can be carried out in many different circumstances and in many different ways.

[8] In the I spy experiment, of course, this judgment is mistaken. But this mistaken judgment need not point to our "natural" tendency to consider our conscious thoughts to be the causes of our actions. It can just as well be explained by noting that subjects judge their movements to be their own because the matching conscious intention just prior to it is, under normal circumstances, indicative of having had an (unconscious) urge to make that move.

References

Anderson, C. J. (2003). The psychology of doing nothing: Forms of decision avoidance result from reason and emotion. *Psychological Bulletin*, *129*(1), 139–166.

Bargh, J. A., Lee-Chai, A., Barndollar, K., Gollwitzer, P. M., & Trötschel, R. (2001). The automated will: Nonconscious activation and pursuit of behavioral goals. *Journal of Personality and Social Psychology*, *81*(6), 1014–1027.

Baumeister, R. F., DeWall, C. N., & Vohs, K. D. (2008). Is the conscious self a help, a hindrance, or an irrelevance to the creative process? *Advances in Psychology Research*, *53*, 137–152.

Baumeister, R. F., & Masicampo, E. J. (2010). Conscious thought is for facilitating social and cultural interactions: How mental simulations serve the animal-culture interface. *Psychological Review*, *117*(3), 945–971.

Baumeister, R. F., Masicampo, E. J., & Vohs, K. D. (2011). Do conscious thoughts cause behavior? *Annual Review of Psychology*, *62*(1), 331–361.

Bayne, T. (2006). Phenomenology and the feeling of doing: Wegner on the conscious will. In S. Pockett, W. P. Banks, & S. Gallagher (Eds.), *Does consciousness cause behavior?* (pp. 169–185). Cambridge, MA: MIT Press.

Bayne, T. (2008). The phenomonology of agency. *Philosophy Compass: Mind and Cognitive Science*, *3*(1), 182–202.

Beattie, J., Baron, J., Hershey, J. C., & Spranca, M. D. (1994). Psychological determinants of decision attitude. *Journal of Behavioral Decision Making*, *7*(2), 129–144.

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, *18*, 227–287.

Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences*, *9*(2), 46–52.

Brand, M. (1984). *Intending and acting: Toward a naturalized action theory.* Cambridge, MA: MIT Press.

Brass, M., & Haggard, P. (2008). The what, when, whether model of intentional action. *The Neuroscientist*, *14*(4), 319–325.

Dijksterhuis, A., Bos, M. W., Nordgren, L. F., & van Baaren, R. B. (2006). On making the right choice: The deliberation-without-attention effect. *Science*, *311*(5763), 1005–1007.

Dijksterhuis, A., Chartrand, T. L., & Aarts, H. (2007). Effects of priming and perception on social behavior and goal pursuit. In J. A. Bargh (Ed.), *Social psychology and the unconscious: The automaticity of higher mental processes* (pp. 51–132). Philadelphia: Psychology Press.

Dijksterhuis, A., & Nordgren, L. F. (2006). A theory of unconscious thought. *Perspectives on Psychological Science*, *1*(2), 95–109.

Dretske, F. (1988). *Explaining behavior: Reasons in a world of causes*. Cambridge, MA: MIT Press.

Frankfurt, H. G. (1971). Freedom of the will and the concept of a person. *The Journal of Philosophy*, *68*(1), 5–20.

Frankfurt, H. G. (1988). *The importance of what we care about: Philosophical essays*. Cambridge: Cambridge University Press.

Fried, I., van Baaren, R. B., & Kreiman, G. (2011). Internally generated preactivation of single neurons in human medial frontal cortex predicts volition. *Neuron*, *69*(3), 548–562.

Gazzaniga, M. S. (1988). Brain modularity: Towards a philosophy of conscious experience. In A. J. Marcel & E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 218–238). Oxford: Clarendon.

Gazzaniga, M. S. (1995). Consciousness and the cerebral hemispheres. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 1391–1400). Cambridge, MA: MIT Press.

Gollwitzer, P. M. (1999). Implementation intentions: Strong effects of simple plans. *American Psychologist*, *54*(7), 493–503.

Gollwitzer, P. M., & Sheeran, P. (2006). Implementation intentions and goal achievement: A meta-analysis of effects and processes. In M. P. Zanna (Ed.), *Advances in experimental social psychology, Vol. 38* (pp. 69–119). New York: Academic Press.

Haggard, P., & Eimer, M. (1999). On the relation between brain potentials and conscious awareness. *Experimental Brain Research*, *126*, 128–133.

Hassin, R. R., Uleman, J. S. & Bargh, J. A. (Eds.). (2004). *The new unconscious* (1st ed.). Oxford: Oxford University Press.

Heil, J. & Mele, A. R. (Eds.). (1993). *Mental causation*. Oxford: Oxford University Press.

Horgan, T., & Tienson, J. (2002). The intentionality of phenomenology and the phenomenology of intentionality. In D. J. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings* (pp. 520–533). New York: Oxford University Press.

Horgan, T., Tienson, J., & Graham, S. (2003). The phenomenology of first-person agency. In S. Walter & H. Heckmann (Eds.), *Physicalism and mental causation: The metaphysics of mind and action* (pp. 323–340). Charlottesville, VA: Imprint Academic.

James, W. (2001). *Psychology: The briefer course*. Mineola, NY: Dover.

Janis, I. L., & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice, and commitment*. New York: Free Press.

Kim, J. (1998). *Mind in a physical world: An essay on the mind-body problem and mental causation*. Cambridge, MA: MIT Press.

Kim, J. (2005). *Physicalism, or something near enough*. Princeton, NJ: Princeton University Press.
Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sceinces*, *8*, 529–566.
Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain*, *106*(3), 623–642.
Mele, A. R. (1992). *Irrationality*. New York: Oxford University Press.
Mele, A. R. (2006). Free will: Theories, analysis, and data. In S. Pockett, W. P. Banks, & S. Gallagher (Eds.), *Does consciousness cause behavior?* (pp. 187–205). Cambridge, MA: MIT Press.
Mele, A. R. (2009). *Effective intentions: The power of conscious will*. Oxford: Oxford University Press.
Morsella, E. (2005). The function of phenomenal states: Supramodular interaction theory. *Psychological Review*, *112*(4), 1000–1021.
Nahmias, E. (2010). Scientific challenges to free will. In T. O'Connor & C. Sandis (Eds.), *A companion to the philosophy of action* (pp. 345–356). Malden, MA: Wiley-Blackwell.
Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, *84*, 231–269.
Pacherie, E. (2006). Towards a dynamic theory of intentions. In S. Pocket, W. P. Banks, & S. Gallagher (Eds.), *Does consciousness cause behavior? An investigation of the nature of volition* (pp. 145–167). Cambridge, MA: MIT Press.
Papies, E. K., Aarts, H., & De Vries, N. K. (2009). Planning is for doing. Implementation intentions go beyond the mere creation of goal-directed associations. *Journal of Experimental Social Psychology*, *45*, 1148–1151.
Richard, R., Van Der Pligt, J., & De Vries, N. (1996). Anticipated regret and time perspective: Changing sexual risk-taking behavior. *Journal of Behavioral Decision Making*, *9*(3), 185–199.
Richardson, R. D. (2007). *William James: In the maelstrom of American modernism*. Boston: Houghton Mifflin Harcourt.
Roessler, J., & Eilan, N. (2003). *Agency and self-awareness: Issues in philosophy and psychology*. New York: Oxford University Press.
Ryle, G. (1949). *The concept of mind*. London: Hutchinson House.
Schurger, A., Sitt, J. D., & Dehaene, S. (2012). An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(42), E2904–E2913.
Simonson, I. (1992). The influence of anticipating regret and responsibility on purchase decisions. *Journal of Consumer Research*, *19*, 105–118.
Soon, C. S., Brass, M., Heinze, H.-J., & Haynes, J.-D. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, *11*(5), 543–545.
Tetlock, P. E., & Boettger, R. (1994). Accountability amplifies the status quo effect when change creates victims. *Journal of Behavioral Decision Making*, *7*(1), 1–23.
Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.
Wegner, D. M. (2005). Who is the controller of controlled processes? In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 19–58). Oxford: Oxford University Press.
Wegner, D. M., & Wheatley, T. (1999). Apparent mental causation. *American Psychologist*, *25*(7), 480–492.
Wilson, T. D. (2002). *Strangers to ourselves: Discovering the adaptive unconscious*. Cambridge, MA: Harvard University Press.