

Introduction: Free Will, Consciousness, and Intentional Action

One of the most pressing questions in daily life, philosophy, and recently also in psychology and neuroscience is whether what we decide and do is in some sense *up to us*. If I take a plane to the other side of the world to work there, marry someone, or have children, or decide to do neither of these things, are these really *my* decisions and things *I* do?¹ Most people probably assume that they can make these decisions freely and that it is up to them what they end up doing. Similarly, we feel it is up to us whether we flex our wrist at a certain moment, decide to interrupt someone or not, push the left or the right button, or decide that one candidate is more suitable to be police chief than the other one. However, arguing for whether and why this is true is notoriously hard.

Traditionally in philosophy, the questions whether we have free will and which conditions need to hold for us to have it are discussed in relation to the truth (or falsity) of determinism. If determinism is true, a complete description of the condition of the entire universe at any point in time combined with a complete statement of the laws of nature entails a complete description of the condition of the entire universe at any other point in time (Mele, 2009, p. 150). The question is whether we have free will even if the condition of the universe *after* you end up doing one thing rather than the other is determined by the condition of the universe as it was at some point, for example before you were born, and the laws of nature. There seems to be a problem for free will here, because then every decision and everything I do, that I am typing this introduction at the moment, what job I am going to have next, and who I will spend my life with, was in some sense already settled before I could have had an influence on what I do. Some take this to be the end of free will, because it suggests that my decisions and deliberations about these important questions do not make a difference to what I do. This is what incompatibilists, i.e., hard determinists and libertarians, agree on: if determinism is true free will cannot exist. They think that free will is about the ability to do otherwise - both me

¹ With 'what I/we do' and 'doings' I include all behavior and bodily movements that are in some sense things we bring about. Intentional actions, unintentional actions, and behavior all fall under this broad category. I explain how I distinguish between these categories in section 3 of this chapter.

having the job I actually end up having and having a different one have to be possible futures for free will to exist - and that if determinism is true we do not have this ability. Libertarians defend the position that free will exists and that determinism is false, while hard determinists think that determinism is true and therefore free will does not exist.

However, most philosophers think that determinism and free will are *compatible* (Bourget & Chalmers, 2014): we can have free will even if determinism is true. One reason for defending this position is that libertarians face a problem: if *indeterminism* is true and if that is necessary for us to have free will, how should we understand deciding to do one thing rather than the other as being up to me? If it is undetermined whether I will do A rather than B, can we then say that me ending up doing A is *my decision*? Imagine that I have to choose between visiting my family or going to a concert, and the condition of the universe, including my reasons, is exactly the same right before I end up choosing either option. Then, what makes me do one thing rather than the other? Again, it seems that my reasons for visiting my family or going to the concert do not make a difference. In general, compatibilists think that free will is about acting for reasons and that we have this ability independently of the truth of determinism. Even if what I decide to do is the result of the condition of the entire universe at some point and the laws of nature, it does not mean that my reasoning abilities are *bypassed*. My reasons to choose one job over another still play a role in what I end up doing. The suggestion is that a distinction should be made between determinism on the one hand and compulsion and coercion on the other. While in the latter cases the agent's ability to act for (the right) reasons is diminished, determinism does not affect this ability. According to many compatibilists there is some sense in which I could have done otherwise, even if determinism is true. I still could have made a different decision if there was reason to do so.

Although this debate between hard determinists, compatibilists, and libertarians is still very much alive (see, e.g., Kane, 2011; Levy & McKenna, 2009; Van Inwagen, 2000), this discussion will not be the main focus of this book. Even though it will become clear later in the book that this discussion does play a role, I will mainly concentrate on the recent contributions of neuroscientific and psychological research to the debate on free will. Neuroscientific studies seem to show that what we do is not the result of conscious states and processes, but instead of unconscious neural activity. Even if we feel we freely choose what to do, research seems to show that in fact the conscious intention 'bubbles up' from our brains and we only become aware of a decision that has already been made (e.g., Haggard & Eimer, 1999; Libet, 1985; Soon, Brass, Heinze & Haynes, 2008). The suggestion is that every time we seemingly freely act the conscious intention is too late to play a causal role, and what we do is the result of unconscious neural activity (see, e.g., Schlosser, 2014; Wegner, 2002 for this interpretation of the findings). Related, psychological research seems to show that conscious states and processes play a much less prominent role in what we do than we might have expected (for an overview see, e.g., Bargh & Chartrand, 1999; Bargh & Ferguson, 2000; Caruso, 2012; Custers & Aarts, 2010; Dijksterhuis & Aarts, 2010). The findings suggest that what agents do is more often than we might have thought the result of external triggers that activate unconscious

states, for example stereotypes, and processes that directly cause us to do things, without conscious states and processes mediating or agents being aware of the causes of what they do. In contrast to the neuroscientific studies, these experiments have not been taken to show that conscious states and processes are never causally efficacious, but their results seem to suggest that very often what we do can be explained solely in terms of external triggers and unconscious states and processes.

Both lines of research seem to undermine one central aspect of free will. If free will is about what I do and decide being up to me, it seems that conscious states and processes have to play an important role. However, if these studies show that they do not, it suggests that indeed it is not up to us whether we flex our wrist at a certain moment, decide to interrupt someone or not, push the left or the right button, or decide that one candidate is more suitable to be a police chief than the other one. If it is really true that conscious states and processes play no causal role at all or hardly do, this suggests that what we do is never or much less often than we thought up to us and free will does not exist or is much more limited than we might have expected. Furthermore, what makes these findings even more disturbing is that we often seem to *experience* that we have freely decided what to do. Therefore, these findings also seem to show that the perspective of the acting agent is unreliable. We feel that we freely decide and consciously intend what to do and that this has an influence on what we end up doing. But if these experiments really show that this is not the case, our experience of consciously deciding and acting is not in line with how our behavior actually comes about. Hence, this experience is an illusion.

In this book I will take up this issue by examining the assumptions about free will, intentional action, and conscious states that play a central role in the arguments against the existence of free will based on this research. Also, I investigate what these studies precisely show about the causal role of conscious states and processes in what we do. Although I will argue that the causal contribution of conscious intentions to what we do is necessary for free will to exist, I conclude that the experimental findings do not show that they do not play this causal role or hardly do. In chapter 2 and 3 I argue that causally efficacious conscious intentions are necessary for intentional action and, as I will explain later in this introduction, for free will. In chapter 4 to 6 I argue that neuroscientific and psychological experiments do not show that explanations of what we do in terms of unconscious neural processes, or external triggers and unconscious states and processes replace explanations in terms of conscious intentions.

In this chapter I will set the stage for this line of argumentation. My main aim is to argue that the reason that the neuroscientific and psychological experiments are, or should be, seen as a threat to free will is not that they show that determinism is true. Also, the threat is not that these experiments show that free will does not exist because our conscious states are not uncaused causes of what we do. This interpretation rests on a problematic definition of free will. Instead, the experiments are threatening because they might show that conscious intentions do not, or less often than we might think, cause what we do. And if that is true, the conclusion is that we do not or only rarely act intentionally, and therefore that free will does

not exist or is severely limited.

I start by explaining the main experiments and interpretations that are taken to undermine free will in section 1 and 2. I will not take into account all scientific threats to free will, e.g., I leave out those that focus on (deterministic) relationships between genes, brain development, and behavior (Bókkon, Vas, Császár & Lukács, 2013; Cashmore, 2010; Lamme, 2011; Swaab, 2010). Rather, I focus on those threats that suggest that in concrete cases what we do is the result of unconscious (neural) states and processes and external triggers rather than conscious states and processes. In the final section of this introduction, section 4, I explain how the five remaining chapters of this book contribute to my two main aims: arguing that causally efficacious conscious intentions are necessary for intentional action, and that the neuroscientific and psychological experiments do not show that explanations of what we do in terms of unconscious neural processes, or external triggers and unconscious states and processes replace explanations in terms of conscious intentions.

1. Neuroscience and Free Will: The Brain ‘Decides’

Benjamin Libet’s (e.g., 1985) experiments in the early 80s have stirred up the debate on free will. What he aimed to do was to combine the results from earlier research (Kornhuber & Deecke, 1965) that showed that voluntary, spontaneous actions are preceded by preparatory activity in the brain, the readiness potential (RP), with measurement of conscious awareness of the intention (or urge) to perform the action. Libet’s main question was how the neural activity, the conscious intention to perform the action, and the bodily movement are related in case of voluntary action. In his experiments Libet (1985) used the following definition of voluntary action:

- (a) it arises endogenously, not in direct response to an external stimulus or cue;
 - (b) there are no externally imposed restrictions or compulsions that directly or immediately control subjects’ initiation and performance of the act; and (c)
 - most important, subjects *feel* introspectively that they are performing the act on their own initiative and that they *are free* to start or not to start the act as they wish.
- (pp.529-530)

The subjects’ task in the experiments was to freely and spontaneously flex their wrist or push a button, whenever they felt the urge to do so. Subjects were explicitly instructed not to plan their movements, but to conduct the action spontaneously and not pay attention to the ‘prospect’ of acting in advance (Libet, 1985, p. 530). This means that the subjects were free to choose to perform this act at any time the desire, urge, decision, or will arises in them.

In the experiments, Libet measured (1) brain activity prior to the performance of the voluntary action (RP), (2) the moment at which the subject had the conscious urge or intention to move (‘W’), and (3) the onset of the bodily movement. The RP is “a scalp-recorded slow

negative shift in electrical potential generated by the brain and beginning up to a second or more before a self-paced, apparently voluntary act” (Libet, 1985, p. 529). ‘W’ can be interpreted as the experience of free will; the moment we have a conscious intention, or are conscious of the intention, to make the movement. In measuring W, Libet asked the subjects to report the moment when they became aware of the wish or urge to act. To achieve this, Libet used a clock with a dot moving in a circle. The subjects had to remember the location of the dot on the clock when they first became aware of the intention to move and report this later.

What made Libet’s experiments the starting point of a new kind of free will skepticism is that the data showed that the subjects became conscious of the urge or intention to move their wrists or finger about 350ms *after* the RP has started and 200ms before the movement. The neural buildup turned out to start already before the subjects had a conscious urge or intention to move, which suggests that the subjects do not consciously decide to move, but only become aware of a ‘decision’ that has already been made by the brain. We might *feel* that we freely decide when to make this movement, but this is an illusion we have only because this preceding neural process occurs outside of our conscious awareness. Libet concludes that “[t]he brain ‘decides’ to initiate or, at least, to prepare to initiate the act before there is any reportable subjective awareness that such a decision has taken place” (1985, p. 536). Conscious states and processes seem to be too late to cause the bodily movements and this was and still is taken by many to have enormous implications: if conscious states do not cause our voluntary actions and our brain ‘decides,’ how can we be free to decide what we are going to do? Furthermore, our experience is apparently unreliable: we may feel that what we do is up to us, but that is only because we do not realize that there is this neural process that has already ‘decided’ what we are going to do.

Libet (1985) thought that these results could be generalized to all voluntary acts, including those actions that are the result of conscious deliberation. Even if agents deliberate about what to do “[t]he urge or intention actually to perform the voluntary act would then still be initiated unconsciously, regardless of the preceding kinds of deliberative processes” (Libet, 1985, p. 539). However, Libet did not want to conclude that conscious states and processes play no role at all. Conscious awareness might appear after the onset of the RP but since it still occurs before the bodily movement, the agent could still *veto* and decide not to act in the way that has been prepared unconsciously (Libet, 1985, p. 538). In Libet’s (1999, p. 52) words: “The conscious will [...] selects which of these initiatives may go forward to an action or which ones to veto and abort, with no act appearing.” If we are able to consciously veto, we might still have some conscious control over what we do. The suggestion is that we might not have free will, but we do have free *won’t*. However, this would be a limited kind of freedom since we still do not consciously bring about what we do, but can only stop processes that started unconsciously.

Other neuroscientists have criticized Libet’s (1985) experiments and the conclusions that are drawn from them. The possibilities of fMRI and further elaboration on the topic has led them to improve the experiments, address some of the difficulties that were pointed out in

response to Libet's experiments, and study not only the timing of the voluntary action but also 'decisions' about *what* movement to make and *whether* to move (see, e.g., Filevich, Kühn & Haggard, 2013; Haggard & Eimer 1999; Soon et al., 2008). For the purposes of my book these improvements and adaptations do not amount to changes that I need to address in detail, because the definition they hold of voluntary action and the instruction to the subjects was similar: make the instructed bodily movement, or veto it, whenever you feel the urge or intention to do so, and do not preplan the movements. Furthermore, these studies underscore Libet's (1985) findings: 'decisions' about what movement to make and whether to move are also preceded by unconscious neural activity. That means that from these experiments the same conclusion is drawn: it is not the conscious agent that decides when to move, what movement to make, or whether to move, but the brain. These findings suggests that free won't does not exist either, because this decision seems to 'bubble up' from the brain as well.

So should we be worried about what these experiments show? Do we experience ourselves to make free decisions while in fact our brain 'decides' for us? Different aspects of Libet's (1985) studies and other Libet-style experiments have been discussed and criticized by many, focusing on problems with the experimental design, definitions used, and conclusions drawn, and experiments have been conducted to undermine the findings (to name a few studies that provide a different perspective on the experimental results and their interpretation, and papers in which the experiments and/or conclusions are criticized: Fischborn, 2016; Gallagher, 2006; Glas, 2007; Gomes, 1998; Kolk, 2012; Matsushashi & Hallett, 2008; Mele, 2009; Oomen, 2013; Pacherie, 2014; Radder & Meynen, 2012; Roskies, 2010; Schlegel et al., 2013; Schlosser, 2014; Schurger, Sitt, & Dehaene, 2012).² In line with some of these responses to the experiments, in my book I do not focus on methodological problems with the experiments or whether it can be concluded that there are neural precursors of voluntary actions on the basis of these experiments. Rather my starting point is: what if we assume that the experiments are sound and there really is this neural activity that consistently precedes the voluntary actions they study. Does that amount to a serious threat to free will? And why is that the case?

In section 3 of this chapter I argue that these experiments could be seriously threatening. The experiments might show that conscious states and processes are not causally efficacious, and because of that we do not act intentionally. But first I will address two ways in which these experiments have been taken to be threatening to free will and dismiss those. It has been suggested that these experiments show that (1) determinism is true, and (2) that what we do is not caused by an uncaused conscious cause, which is, mistakenly I think, taken to be necessary for free will by some. To start with the first suggestion, the findings of Libet-style experiments have been taken to show that determinism is true, and that unconscious neural activity determines what we are going to do. But this conclusion does not follow from the experiments. For one thing, since only the temporal characteristics of voluntary actions are measured, no

² Later in the book I will engage with the relevant arguments and interpretations that are addressed in some of these papers in more detail.

conclusions about determinism or even causality can be drawn from these experiments.³ Often an event consistently takes place after another event but there is no causal relationship. For example, even though every time I go to bed I take off my glasses right before, that does not mean that taking my glasses off *causes* me to go to bed. The only thing that can be concluded is that there is a correlation between the neural activity, the conscious intention, and the bodily movements. Furthermore, for compatibilists the truth of determinism would not threaten the existence of free will, as I pointed out earlier in this introduction. Thus, the threat of these experiments is not that they show that determinism is true.⁴

Second, neuroscientists that conducted these experiments and draw conclusions about free will from them tend to assume that for free will the conscious agent or state needs to be an uncaused cause, a point of origin of action beyond which a causal chain cannot be traced (Bayne, 2011). As Wegner (2002) states it, “[t]he Free Willer must be a mechanism that is *unresponsive to any past influence*” (p. 322). But this definition is problematic, not only if compared to the definitions of free will in philosophy, but also if we compare it to folk intuitions about free will (e.g., Nahmias, 2006). For example, most philosophers, libertarians and compatibilists alike, accept that agents are responsive to past influences, that reasons and facts about the world play a role in what they do and the choices they make. And even if they think that dualism is true or needs to be for free will to exist, they would not think it is a problem for free will that correlations between mental states and neural states can be found. Also, since most philosophers, as well as neuroscientists, take a naturalistic stance, they reject a notion of free will that depends on an uncaused cause (see, e.g., Bayne, 2011; Mele, 2007a). The difference is that philosophers think that this definition of free will is the problem, while neuroscientists think that the notion of free will precisely *is* about uncaused causation, and conclude that free will does not exist. But, there is good reason to reject this definition of free will. It is not clear at all how being an uncaused cause gives us the kind of control over what we do and the kind of relationship to our bodily movements that we need to exercise control. Why would a decision I make in which I am unresponsive to my environment and my personal history be *my* decision in any interesting sense? And how would such an independent ‘decider’ even affect bodily movements? Thus, also in this way the experiments do not pose a threat to free will, because they do not undermine an interesting definition of free will by taking it to be

³ This points to a problem for the neuroscientific study of free will as neuroscientists often define free will. Since according to them free will needs to be an uncaused cause, studying it does not allow for manipulation. However, at the same time manipulation is necessary to conduct a real experiment and draw causal conclusions (e.g., Dooley, 2001). Therefore studying free will with this definition in place cannot lead to the conclusion that what subjects do is caused by the preceding neural activity.

⁴ One might also interpret the threat of these findings being that what we are going to do can be predicted accurately and that this undermines free will. This is for example central in the experiment by Soon et al. (2008), in which decisions about what movement to make could be predicted well in advance on the basis of preceding neural activity. However, the results are far from predicting what a subject will do with high accuracy. Furthermore, Holton (2013) argues that, even though it would be a problem for free will if we could predict with full accuracy what someone will do, it is unlikely that this level of prediction is possible.

an uncaused cause.

As I will explain in more detail in section 3, a much more threatening interpretation of these findings is that they might undermine the causal efficacy of conscious intentions and show that we do not act intentionally, even though we might think we do. First, however, I will put forth the psychological research that is taken to be threatening to free will.

2. Psychology and Free Will: The ‘Adaptive’ Unconscious

The role of unconscious states and processes in what we do and its relation to free will is a central topic of discussion in some areas of psychology as well. Studies seem to show that conscious states and processes play a much less prominent role in what we do than we might have expected (for an overview see, e.g., Bargh & Chartrand, 1999; Bargh & Ferguson, 2000; Caruso, 2012; Custers & Aarts, 2010; Dijksterhuis & Aarts, 2010), and that unconscious states and processes can trigger and guide higher cognitive acts and decision-making without conscious guidance (e.g., Bargh & Ferguson, 2000). The suggestion is that even in more complex situations, we do not need conscious thought or will to function properly.

In contrast to neuroscience, there is not one psychological experiment that started the discussion about free will in psychology and that has been repeated in several variants. Rather, a whole collection of studies seem to all point to the conclusion that what we do is often the result of external triggers that directly activate unconscious states and processes that cause behavior. A phenomenon that has been studied (and criticized⁵) a lot is priming. In cases of priming, the suggestion is that external triggers activate unconscious goal representations that directly cause and guide what we do (e.g., Custers & Aarts, 2010). Many experiments on priming effects have been conducted. To give a few examples, priming studies show that being primed with stereotypes of elderly makes us walk slower, and priming with professors leads us to perform better on a knowledge quiz (Bargh, 2006).

Furthermore, research shows that personal goals and motivation, for example the goal to perform well or to be egalitarian (Caruso, 2012, p. 124), can be the result of external triggers that directly activate unconscious states and processes. Even when activated unconsciously, this goal can play the same role in what agents do compared to when the subjects consciously decided to act on this goal (Bargh, 2008, pp. 142-143; Bargh & Ferguson, 2000, p. 932). Experiments also show that priming can enhance motivation and improve performance on tasks (e.g., Bargh & Gollwitzer, 1994). Similarly, unconscious responses to stimuli can make us perform better on a task. In an experiment by Bechara, Damasio, Tranel, and Damasio (1997), it was found that subjects began to make better decisions about which cards to pick to win money before they consciously realized which strategy worked best.

⁵ The reproducibility of priming studies has been subject of heated debate (see, e.g., Open Science Collaboration, 2015). However, since the worst-case scenario is that they do exist and we might find more support for it in the future, it is interesting to examine what this might mean for free will. Therefore I will not concern myself with this debate but accept that priming effects exist.

Thus, it seems that what we do can be directly influenced by unconscious states and processes, not only in the case of simple responses like adapting our posture which we might expect, but also in case of complex social behavior and acting on goals (Bargh & Ferguson, 2000, p. 928). The findings seem to show that it does not matter whether goals are activated and operate consciously or unconsciously to have an effect on what we do (Bargh & Ferguson, 2000, p. 936). Wilson (2002, p. 36) sums it up as follows: “Our conscious mind is often too slow to figure out what the best course of action is, so our nonconscious mind does the job for us and sends us signals (e.g., gut feelings) that tell us what to do.”

The experiments I cited above mainly purport to show that unconscious states processes can be adaptive and help us perform better or just slightly different. The lack of involvement of conscious states and processes does not seem to be worrying. However, other experiments seem to show that unconscious influences can also make us do things that seem to be out of line with our explicit attitudes and principles, and that also in those cases what we do is not the result of conscious deliberation and/or a conscious intention to act in a certain way. For example, in an experiment by Darley and Batson (1973) subjects’ willingness to help a stranger in distress was influenced by whether the subjects thought they were in a hurry or not. A different example is an experiment that is conducted in which agents are primed to interrupt (or not interrupt) the researcher. In this experiment (Bargh, Chen & Burrows 1996, pp. 233-234), after working on an assignment that primed subjects with rudeness, politeness, or neither of the two, subjects tended to interrupt the researcher more often when primed with rudeness compared to the other two conditions (Bargh et al., 1996, p. 235). And Schnall, Haidt, Clore, and Jordan (2008) found that, without the subjects being aware of this influence, disgust can increase the severity of moral judgments. Finally, there is the case of implicit bias, in which people’s judgments and what they do is unconsciously influenced by the race, gender, or other characteristics of the person they are dealing with. For example, in an experiment by Uhlmann and Cohen (2005), on selecting applicants for a job as police chief, the experimenters found that the subjects “tailored their criteria to favor whatever qualities the individual applicant of the desired gender happened to have” (p. 479). Since subjects (unconsciously) aimed to select the male candidate, they found the candidate being streetwise the most important criterion if the male candidate possessed this characteristic, but not if the female candidate was most streetwise.

Thus, these experiments seem to show that, even in morally relevant situations, a lot of what we decide and do is not the result of conscious deliberation and conscious intentions, but of unconscious mental states and processes that are triggered by environmental features and that produce behavior outside of conscious awareness and guidance (Bargh & Chartrand, 1999, p. 462). To make the claim more specific, Bargh’s (1997, p. 243) suggestion is that our psychological reactions from moment to moment are 99.44% automatic.

The question is how to precisely interpret these findings, and why they might show that we do not have free will, or that free will is severely limited. Similar to the neuroscientific studies, it is sometimes argued that these experiments show that determinism is true. Also,

the assumption that for free will the conscious agent needs to be an uncaused cause plays a central role. For example in the paper by Bargh and Ferguson (2000), the central claim is that these findings show that not only unconscious processes are deterministic, but that conscious processes are as well. Their conclusion is not so much that conscious states and processes play a marginal role, but rather that there is no substantial difference between unconscious and conscious states and processes because they can and do fulfill the same functions (Bargh & Ferguson, 2000, p. 938). Even though one could conclude from this that, similar to conscious processes, unconscious processes are willful, intentional and free, Bargh and Ferguson draw the opposite conclusion, namely that conscious processes are mechanistic and determined as well. According to Bargh and colleagues it is mistake to think that only automatic processes are caused and have underlying mechanisms, and that controlled processes do not and are therefore “free” (Bargh, 2006, p. 42).

Similar to what I said about the neuroscientific findings, the question is whether appealing to determinism amounts to an interesting way to spell out why these experiments pose a threat to free will. I think it is not. For one thing, the truth of determinism does not follow from the experiments. Instead, it is a working assumption that is held in conducting this kind of research, which Bargh and Ferguson (2000) also acknowledge. Unlike the researchers that conduct Libet-style experiments, these psychologists are in a position to conclude that there is a causal relationship between the external trigger and what the subjects do. On the basis of how the experiments are conducted, in which a comparison is made between subjects that are exposed to the trigger and subjects that are not, it can be concluded that the external trigger actually makes a difference to what the subjects do (see, e.g., Dooley, 2001). However, even if they are in a position to draw this conclusion, this does not show that determinism is true, or that a deterministic relationship between the trigger and the behavior exists. Furthermore, I should point out again that according to most philosophers determinism is not a problem for free will. So claiming that these experiments show that determinism is true and that therefore free will does not exist is not an interesting way to spell out the threat to free will from psychological research.

Furthermore, researchers such as Bargh (e.g., 2008) hold a definition of free will that diverges strongly from the dominant definitions in philosophy. For Bargh (2008) free will “requires an absence of both external and internal determination of the action” (p. 130). This suggests that the conclusion that the psychological experiments show that free will does not exist depends on a notion of free will as an uncaused cause of what we do. And since these experiments are uncovering all kinds of causes of what we do, we no longer have to posit a ‘free willer’ to explain what is taking place. This source of explanation has become redundant: “[i]t is not necessary to invoke the idea of free will or a nondetermined version of consciousness as a causal explanatory mechanism in accounting for higher mental processes in humans” (Bargh & Ferguson, 2000, p. 939). And even though we feel we act freely, this experience is misleading because we are simply not aware of the real causes of what we do, also in the case of conscious processes, and therefore we experience what we do as ‘free’ (Bargh, 2008, p. 131

& p. 133). However, as I pointed out in relation to the neuroscientific findings as well, the idea that free will requires absence of both external and internal determination is problematic.

We can conclude that our experience of free will understood as being an uncaused cause cannot be ‘the illusion of free will’ that is at stake. Not only is it questionable whether our experience of free will amounts to being an uncaused cause (see, e.g., Deery, Bedke & Nichols, 2013; Nahmias, Morris, Nadelhoffer & Turner, 2004 for studies on the phenomenology of free will), since being an uncaused cause is not what free will is about, this experience, even if we have it, cannot amount to the illusory experience of free will.

There is a third way in which the findings from psychological research could be taken to threaten free will. They might show that these influences make us act against our intentions and values. However, this should be treated as a problem for free will apart from the problem of unconscious factors playing a causal role. We can also act against our intentions and values knowingly and consciously, and this might similarly threaten our freedom (see, e.g., Frankfurt, 1971; Wolf, 1990). Thus what is specifically threatening about these findings cannot just be that we act against our intentions and values. We have to distinguish between this issue and the role of unconscious states and processes, or at least focus on why this combination might be particularly threatening.

3. The Threat: Consciousness Is Not in the Right Way Involved in What We Do

Then why are the neuroscientific and psychological experiments regarded as threatening to free will, if they use a problematic definition of free will and do not show that determinism is true? To answer this question we need to assess the relationship between consciousness and free will in more detail; the role of conscious and unconscious processes in what we do is what is at stake. We often feel that we consciously decide what to do and that conscious intentions play an important role in what we end up doing. But what if psychology and neuroscience show that this is an illusion? Even though we experience that our conscious processes play a causal role, in fact unconscious (neural) processes and external triggers might directly cause what we do. Apart from whether determinism is true or whether what we do is the result of conscious uncaused causes this would be threatening to free will, because it seems to show that what we end up doing merely happens to us and we lack conscious control. This intuition is confirmed by experimental philosophy: lay people do think that conscious deliberation and conscious intending are important for free will, in the sense that the conscious agent should not be ‘bypassed’ (see, e.g., Nahmias, 2006). Hence, if these neuroscientific experiments really show that all our ‘voluntary actions’ are caused by unconscious neural processes and they are not the result of conscious states and processes, we would have to be worried whether what we do is ever really up to us.

The threat from psychological experiments is similar, in that these studies seem to show that conscious states and processes only play a minor role in what we do, and more often than

we think unconscious processes and external cues are the direct triggers. These psychological explanations are taken to replace explanations of what we do in terms of conscious states and processes. The suggestion is that “conscious intentions and conscious control are not necessary causal precursors for much of what we do” (Vierkant, Kiverstein, & Clark, 2013, p. 8). The experiments do not show that conscious processes *never* cause what we do, but that they do so less often than we think, and other factors play a more dominant role than we might have expected. Nadelhoffer (2010) takes the neuroscientific findings to be a *global* agential threat, because they leave no room for conscious volition at all, and the psychological findings to be a *local* agential threat, because they show that the domain of our conscious agency and control is constrained and smaller than we thought (p. 178). Hence, the general idea is that “to whatever extent unconscious cognitive states and processes determine our actions, to that extent our actions fail to be free” (Caruso, 2012, p. 99). These neuroscientific and psychological experiments seem to show that unconscious processes determine all or a lot of it.

But how exactly is the threat to free will related to a lack of involvement of conscious processes? We have seen that lay people assume this to be problematic. Caruso (2012, e.g., p. 99) also appeals to intuitions and common sense views of free will. He claims that we simply tend to think that free will depends on conscious processes, and that the burden of proof is on those who argue otherwise. But this does not amount to an *argument* that free will depends on conscious processes; our intuitions and common sense views might be wrong. Furthermore, it is often not made explicit why and in what way conscious processes are necessary for free will. Similarly, it is often merely stated that we lack conscious control, but not much is said about what conscious control is and how it is important for free will. Hence, we need to say more about the kind of conscious processes that need to be involved and in what way they need to be, to assess whether these findings undermine the existence of free will. In the remainder of this section I distinguish three ways in which conscious processes can be taken to be of central importance for free will: (1) conscious deliberation as a necessary condition for free will, (2) being conscious of all the factors that contribute to what you do as a necessary condition for free will, and (3) causally efficacious conscious intentions as a necessary condition for free will. After arguing that the first two are not necessary for free will, I conclude that the focus should be on the third condition.

Why Some Conscious Processes and States Are Not Necessary for Free Will

The first suggestion is that conscious deliberation is a necessary condition for free will. Pacherie (2014) suggests this to be the case when she states that “the worry raised by Libet’s data is that the conscious decision to act is arrived at through unconscious brain processes rather than through conscious deliberative processes” (p. 35). A similar interpretation of the psychological findings can be given: if these studies show that what we do is often the

direct result of external triggers and unconscious states and processes, then these agents did not deliberate and presumably could not have deliberated about what they are doing. The assumption is that only if what I did came about as a result of conscious deliberation it was up to me. But this assumption is problematic. I do not have to consciously deliberate about a lot of things I do on a regular basis - for example having oatmeal for breakfast, driving to work, or going for a run every Saturday - but it does not follow from this that these actions are therefore not free. If that were true, then the first time I decided to go for a run on Saturday and did consciously deliberate about it my running would be a free action, while every other time I went it would not be. This would be a strange conclusion. Even if I put on my running gear automatically when I see it is 10am and time to go, I still act for reasons I endorse and there is no reason to think that I could not have done something else if I had reason to. Thus, the idea that these experiments are threatening to free will because the agents did not consciously deliberate is problematic, because this suggests that we are no longer free once we made up our minds about what to do.

This brings me to the second suggestion: that free will requires being conscious of all the factors that contribute to what you do. As Caruso (2012) puts it: “If conscious intention, control, and guidance of behaviour is a folk-psychological requirement for free will, then it’s easy to see the problem. These individuals are completely unaware of what’s causing them to behave as they do” (p. 119). And in Wegner’s (2004) words: “The mind creates this continuous illusion because it *really doesn’t know* what causes its actions” (p. 654). The experimental findings seem to show that we do not have access to these causal mechanisms and processes that actually produce what we do (e.g., Bargh, 2006; Wegner, 2002, p. 97), and this is why we do have the illusion of being free.

Research shows that indeed a lot of what we do depends on contributions of sub-personal motor mechanisms (see, e.g., Jeannerod, 2003), and philosophers that study the role of consciousness in action acknowledge this (e.g., Eilan & Roessler, 2003; Gallagher, 2006; Pacherie, 2014; Shepherd, 2015). An action as simple as picking up a glass takes a lot of feedback and adapting that we (luckily) are not aware of. However, as Gallagher (2006) argues, we should not understand free will in terms of these motor mechanisms. The relevant question is whether I want to pick up the glass and whether I have reasons to do so. I could very well be freely picking up the glass even if I do not have insight into all the specifics of what I do.

Furthermore, in what way I specifically do what I do does not undermine free will. How I precisely walk to the supermarket - the speed, where I cross the road, where I exactly place my foot - is probably influenced by all kinds of factors I am not aware of: the weather, the mood I am in, the level of adrenaline in my blood, maybe even traffic I automatically respond to. Without being aware of it, the smell of freshly baked bread might have made me go to the supermarket now instead of 7 minutes later. But we should not conclude from this, that since I am not aware of all these kinds of influences, my freedom, at least regarding walking to the supermarket, is diminished. I still have reason to walk to the supermarket and I am

committed to doing that, it is in line with my principles, and there is no reason to think that I could not choose to do something else or go later if I have reason to do so.

That means that free will does not require us to be conscious of all the factors that contribute to what we do. However, I do not want to conclude that being aware of what influences you to do what you do is completely irrelevant, because sometimes this kind of unawareness might interfere with abilities that are of central importance for free will. For example Nahmias (2007, p. 179) states that “[t]he less accurate your knowledge of which factors might influence you, the less you can control the influence of those factors you want to make a difference - that is, the less you can act on your principles.” Furthermore, studies show that we sometimes confabulate stories about why we do what we do, because we do not have insight into what actually made us do it (see, e.g., Nisbett & Wilson, 1977). Although this research is often not directly discussed in connection to free will as Libet-style experiments and priming studies are, some take it to be a challenge for (certain views on) free will (see Caruso, 2012, p. 131; Gazzaniga, 2011; Sie & Wouters, 2010). Since, as I will explain in the next section, our ability to act for reasons is often taken to be necessary for free will, this makes perfect sense. However, this also suggests that being conscious of certain influences is more important than being conscious of others, and that this depends on whether it diminishes our ability to act for reasons. Hence, this distinction cannot be made independently of what it means to act for reasons and to act intentionally. This implies that we have to understand this ability and how it relates to free will first.

To conclude, this second necessary condition can be discarded as well, because we do not need to be conscious of all the factors that contribute to what we do to have free will as Bargh (e.g., 2008), Caruso (2012), and Wegner (2002) suggest. It might matter for free will whether we are conscious of certain influences on what we do, but my suggestion is that this can be best understood in relation to what it means to act for reasons and to act intentionally. This brings me to the third necessary condition for free will: the causal role of conscious intentions.

Why Causally Efficacious Conscious Intentions Are Necessary for Free Will

I already pointed out that Libet-style experiments have been interpreted as showing that conscious intentions are too late to cause the bodily movement, and that because of that they are taken to show that we do not have free will. According to Wegner (2002) for example, Libet-style experiments and other studies show that conscious intentions are not causally efficacious but only function as previews of what we are going to do.⁶ Conscious intentions and

⁶ One could argue whether his theory is really based on Libet’s findings or whether it is the result of Wegner’s position on causation and his adherence to epiphenomenalism in light of which he interprets the research findings. Nahmias (2002, p. 53) states that Wegner defends modular epiphenomenalism (rather than ‘complete’ epiphenomenalism) in which only conscious intentions are causally inefficacious, but other conscious states or processes might play a causal role. Although for now I interpret Wegner’s theory in line with Nahmias’ view, throughout the book it will become clear that overall metaphysical assumptions play a large role in conclusions that are drawn from the experiments about free will, the (implicit) assumption that epiphenomenalism is true being one of them.

what agents do are the end points of two separate causal pathways that start unconsciously. That means that there are two *actual* causal paths in Wegner's theory: between unconscious cause of thought and thought and between unconscious cause of action and action. The presumed causal path between thought and action is an *apparent* causal path. We may feel that our conscious intentions cause our bodily movements, if the intention occurs before them (priority), is consistent with them (consistency), and is not accompanied by other potential causes (exclusivity) (Wegner, 2002, p. 69). However, causation is not a property of the conscious intention (Wegner, 2004, p. 652). We merely interpret intentions to be causes, because the intention and what we do are two sequential, consistent, and exclusive events. In fact, conscious processes are too slow to trigger what we do and follow instead of lead. Similarly, as Bargh (2006, p. 38) puts it in relation to the findings from psychological studies: "the individual's behavior is being 'controlled' by external stimuli, not his or her own consciously-accessible intentions or acts of will." Thus, these findings are taken to show that conscious intentions do not cause what we do, even though we feel they do.

Now we have identified a serious threat to free will. Many philosophers think (1) that intentional action is necessary for free will, and (2) that causally efficacious (conscious) intentions are necessary for intentional action.⁷ For example Mele (2009, pp. 154-155) says that "[t]he basic compatibilist idea is (roughly) that when mentally healthy people act intentionally and rationally in the absence of compulsion and coercion they act freely" and "some libertarians maintain that when mentally healthy people act intentionally in the absence of compulsion and coercion they act freely" (Mele, 2009, p. 155). Furthermore, from what most philosophers take intentional actions to be, it is clear that there is a strong connection between intentional action and free will. First, intentional actions are things agents do for reasons (see, e.g., Anscombe, 1957; Davidson, 1963; Goldman, 1970; Mele, 1992) and are committed to or settled on performing (e.g., Clarke, 2010; Mele, 1992). Second, intentional actions are taken to be those events that are controlled by the agent: the agent sets out and is able⁸ to adapt her bodily movements in order to bring about what she intends to bring about. She guides and sustains her behavior towards an aim. In contrast to behavior like eye blinking (in most cases), accidents like falling, and *unintentional* actions, for example waking up the neighbor while redecorating your living room, intentional actions are events that reflect the agent's reasons and commitments and that she controls. Because of this, it makes perfect sense to limit free will to those things agents do intentionally.

Since the issue in these experiments is whether conscious intentions cause bodily

⁷ 'Conscious' is between parentheses because not all philosophers think that intentions have to be conscious for agents to act intentionally (see, e.g., Mele, 2009, chapter 2). In chapter 3 of this book I argue that I am not convinced by Mele's arguments, and that I take it to be the case that intentions that cause intentional actions are necessarily conscious.

⁸ Often, the discussion of whether some event counts as an intentional action depends on where to draw the line between what agents intentionally do and whether what they bring about is due to luck (see, e.g., Mele & Moser, 1994; Shepherd, 2014). How much control is needed for an achieved result to count as intentional? And how close should the actual execution of the action be to the plan for it to count as intentional? Since these considerations are not of main relevance in this book, I do not take them into account.

movements, I am going to take the action theory that takes this causal relationship to be the defining criterion of intentional action as starting point of my analysis in this book. Many philosophers defend this *causal theory of action*. Also, it seems to be (implicitly) adopted by neuroscientists and psychologists, at least in the sense that conscious intentions need to be causally efficacious for what the agent does to be attributed to her and to possibly count as free. According to the causal theory of action, only those bodily movements that are caused in the right way⁹ by an (conscious) intention count as an intentional action.¹⁰ On the basis of the content of the intention that triggers, guides, and sustains the bodily movements, we can distinguish between intentional actions, unintentional actions, and behavior. More specifically, if behaviors are not caused by an intention, they are mere bodily movements rather than actions. If what the agent brings about is not caused by an intention with that content, for example waking up the neighbor while moving the couch to the other side of the room, it is an unintentional action. This means that not all accurate descriptions of what an agent is in fact doing – moving the couch to the other side of the room, waking up the neighbor, moving dust around, scaring the cats, scratching the floor - are also correct descriptions of what she is doing intentionally. Since she had the intention to move the couch to the other side of the room in this scenario and that intention triggered, guided, and sustained what she was doing, she only intentionally did that. Then, according to the causal theory of action the intention both causes and rationalizes the intentional action.

The causal role of conscious intentions and whether we act intentionally should be distinguished from the two other possible relations between consciousness and free will I addressed earlier. Even though neuroscientists and psychologists seem to assume that conscious intentions and intentional actions are necessarily the result of conscious deliberation (see, e.g., Mele, 2009, p. 29 on the commitments of Haggard & Clark, 2003), this is not true according to the causal theory of action. For example, in the case of routine actions like feeding the cats in the morning, I do not have to consciously form this intention every morning and make up my mind about what to do. I just come to have this intention when I, for example, enter the kitchen. However, even if I do not deliberate I do have this intention

⁹ Deviant causal chains are subject of much controversy in action theory. Sometimes reasons or intentions cause ‘actions’ in the wrong way, and the question is whether the causal theory of action has the means to distinguish between deviant causal chains and the right kind of causation. Davidson’s (1973) example: “A climber might want to rid himself of the weight and danger of holding another man on a rope, and he might know that by loosening his hold on the rope he could rid himself of the weight and danger. This belief and want might so unnerve him as to cause him to loosen his hold, and yet it might be the case that he never chose to loosen his hold, nor did he do it intentionally” (pp. 153-154). In this example the climber’s reason (or intention) did cause the behavior, but we do not want to say it was an intentional action. Causalists have put forth several ways to deal with this problem, for example by emphasizing the guiding and sustaining role of intentions (Mele, 1992).

¹⁰ Some philosophers defend a causal theory of action that does not include a role for intentions, but that claims that intentional actions are caused by other psychological states, for example belief-desire pairs (see, e.g., Davidson, 1963). However, most philosophers agree that intentions play a unique functional role and need to be included in the theory (e.g., Brand, 1984; Bratman, 1987; Mele, 1992; Searle, 1983) In this book I focus on the version of the causal theory of action that includes this role for intentions.

that causes me to feed the cats, and I still act for reasons and am committed to doing that. Also, even if I was not conscious of how this intention precisely came about, whether it was the bird singing outside or fluctuations in brain activity, and that I do it now instead of two minutes later and with my left hand instead of my right, I still intentionally feed the cats. That is, independently of whether I deliberate or am aware of all the causal contributions to my intention and intentional action, I can still act intentionally. And what I do intentionally is in some important sense up to me, because I do it for reasons, I am committed to doing it, and I control my behavior in light of the intention I have.

To conclude, if research shows that (conscious) intentions are never, or hardly ever, causally efficacious, we would never act intentionally or much less often than we might have thought. Most philosophers think that this conclusion is threatening the existence of free will, because it would mean that (almost) everything we do ends up being events that happen to us rather than things we do. Furthermore, this also means that we would be under the illusion that we act intentionally while we do not. In line with Wegner's theory (2002), it would show that we might experience that our conscious intentions cause what we do, but in fact they are endpoints of a causal path. Then, free will would be an illusion in the sense that we might feel that we act intentionally but in fact we do not, and therefore we do not act freely either.

4. Preview of the Book

From my analysis in this chapter, I conclude that if we want to assess whether neuroscientific and psychological experiments are threatening to free will, we need to find out whether these experiments show that we never or hardly ever act intentionally. In order to do so, the first step is to settle whether intentional action indeed depends on causally efficacious conscious intentions. The second step is to assess what these experiments precisely show about the causal role of conscious intentions. Do the neuroscientific studies really show that unconscious neural activity causes what subjects do, instead of the conscious intentions subjects do have? And can it be concluded from the psychological experiments that to a large extent our behavior can be explained in terms of external triggers that activate unconscious states and processes that together lead us to do what we do, without any role for conscious intentions? This brings me to the two research questions that are central in this book:

- (1) Are causally efficacious conscious intentions necessary for intentional action?
- (2) Do neuroscientific and psychological experiments show that explanations of what we do in terms of (a) unconscious neural processes, or (b) external triggers and unconscious states and processes replace explanations in terms of conscious intentions?

In chapter 2 and 3 of this book I focus on question 1, whether intentional action depends on causally efficacious conscious intentions, and my conclusion is that it does. In **chapter 2** I first examine whether we can make sense of intentional action without (a causal role of) intentions. I assess an alternative action theory, Frankfurt's (1978) guidance theory,

according to which whether an agent acts intentionally does not depend on an actual causal relationship between the agent, or the intention the agent has, and the bodily movements (see Di Nucci, 2008, 2011, 2013; Pollard, 2006; Zhu, 2004 for a recent defense of this approach to intentional action). I first address Di Nucci's version of the theory, the *guidance without intentions* theory. I argue that without a role for intentions the theory fails to make the right distinctions between intentional actions, unintentional actions, and behavior. Furthermore, it is unclear how agents can come to know about what they are intentionally doing without having an intention in doing so, while this knowledge does play a crucial role in the theory. Because of this, I address a version of the theory that does allow for the agent to have an intention that is not necessarily causally efficacious. However, I argue that no convincing examples have been given by proponents of the guidance view of a case of intentional action in which an agent has an intention that is not causally efficacious. I conclude that intentional action depends on causally efficacious intentions.

In **chapter 3** I examine whether this *occurrent* intention that triggers, guides, and sustains the intentional action necessarily is a *conscious* intention. Neuroscientific and psychological experiments leave open the possibility that subjects act intentionally, because what they do might be caused by an unconscious intention. Most researchers take intentional action and intentions to be conscious by definition, and even though their research shows that unconscious processes are smart, flexible, and automatically controlled (e.g., Bargh & Ferguson, 2000, p. 939), they do not conclude that these unconscious processes might involve unconscious intentions and that agents might still act intentionally and freely. Conversely, Marcel (2003) and Mele (2009) propose that intentions might do the same causal work without the agent being conscious of them. However, I argue that this proposal is not convincing, and claim instead that those intentions that trigger, guide, and sustain intentional actions are best conceived as conscious intentions. I give two reasons for this. First, this would mean that we have to accept that agents can act intentionally without being conscious of what they are doing. This goes against the common conception that intentional actions are things agents are committed to doing. Second, it means that sometimes no one is in a position to determine which accurate descriptions of what the agent does are also correct descriptions of what she does intentionally. I provide alternative explanations for the examples of unconscious occurrent intentions that Marcel and Mele give on the basis of Papineau's (2015) understanding of basic action. I argue that either agents consciously intend to perform this specific action, or they consciously intend to perform an action on a higher level of description, but that this depends on the situation and the skills of the agent.

This means that in the remainder of the book I have to assess what exactly can be concluded from the neuroscientific and psychological experiments about the causal role of conscious intentions and intentional action. This brings me to question 2: Do neuroscientific and psychological experiments show that explanations of what we do in terms of (a) unconscious neural processes, or (b) external triggers and unconscious states and processes replace explanations in terms of conscious intentions? If intentional action depends on

causally efficacious conscious intentions, the suggestion is that we might feel that we act intentionally, but in fact there is a different explanation of what we do. I will take up this question in chapters 4 to 6.

In **chapter 4** I assess what precisely happens in Libet-style experiments and why it is concluded from these experiments that the brain ‘decides’ instead of the conscious agent. I specifically focus on experiments investigating the so-called *free won’t* and argue that definitions and operationalizations of voluntariness and freedom play an important and problematic role in the conclusion that is drawn on the basis of these experiments. The experiments are designed in such a way that the conclusion that the brain ‘decides’ is almost unavoidable, because potential reasons to act on are excluded from the experimental setting. I conclude that, for neuroscientific experiments to provide a valuable contribution to the debate on free will, researchers should at least acknowledge that intentional action and acting for reasons are central to free will. This means that the conclusion that is drawn from these experiments, that unconscious neural processes cause what subjects do, should be rejected.

Given that the conclusion that what subjects do is caused by unconscious neural processes is based on problematic experiments, an interesting question is whether these experiments might still allow for the conclusion that what we do is caused by conscious intentions. I take up this issue in **chapter 5**. The skeptical line of argument on the basis of Libet-style experiments that I will address consists of the following three claims: (a) Libet-style experiments show that the conscious intentions that are the object of study are not the causes of what agents do, (b) what agents do is only free if it is caused by a conscious intention, and (c) the results of Libet-style experiments can be generalized to all conscious intentions. In chapter 5 I argue that the common response to claim (c), pointing out the distinction between proximal and distal intentions, does not capture the distinction that is crucial, which is the one between passively acquired and consciously formed intentions. I argue that if an intention is *consciously formed*, conscious processes necessarily causally contributed to what the agent does, because the intention could not have already been there in unconscious form. Intentions are only consciously formed if the agent is uncertain about what to do and has to make up her mind, i.e., to integrate information in order to arrive at a decision. I show that these conditions are not met in Libet-style experiments. On the basis of my analysis I argue that not only claim (c) but also claim (a) should be discarded, because otherwise we end up with the problematic conclusion that only if we have to make up our minds about what to do, and consciously form an intention, we act intentionally. Rather, it makes more sense to conclude that passively acquired conscious intentions are causally efficacious as well. The experiments do not show that the agent’s experience that conscious intentions are causally efficacious is illusory.

Finally, in **chapter 6** I focus on the psychological research from which it is concluded that causal explanations of what we do in terms of external triggers and unconscious states and processes replace explanations in terms of conscious intentions, and that because of that we lack conscious control. I discuss three prominent experiments on priming and implicit bias to see whether this conclusion follows, and I argue that it does not. I show that often the

findings in these experiments do not exclude the possibility that agents act intentionally as well. Furthermore, I conclude that 'lack of conscious control' is not an interesting notion if not understood in relation to intentions, values, and principles an agent has. I suggest that in every situation in which we intentionally act there are also a lot of accurate descriptions under which we lack this kind of conscious control, but often this lack of control is simply irrelevant. I conclude that there is no reason to think that these experiments show that explanations of what we do in terms of conscious intentions are replaced. In line with our experience as acting agents, there is good reason to think that our conscious intentions often cause what we do.

On the basis of my analysis I conclude (1) that causally efficacious conscious intentions *are* necessary for intentional action, and (2) that neuroscientific and psychological experiments *do not* show that explanations of what we do in terms of (a) unconscious neural processes, or (b) external triggers and unconscious states and processes replace explanations in terms of conscious intentions. We should not conclude from these studies that what we do merely happens to us. Because of this, the experiments do not give us reason to think that our experience as acting agents is unreliable, and that free will is an illusion. In the **epilogue** I say more about the role of science, philosophy, and the perspective of the agent in understanding intentional action and free will. I suggest that science can contribute to our knowledge of intentional action and free will, but that, in order to conduct sound experiments and ask valuable research questions, what intentional actions are and what free will is needs to be established first. My suggestion is that in order to achieve this, philosophy and the perspective of the agent should play an important role.