Chapter 16
BCI and a User's Judgment of Agency

Rutger Vlek, Jan-Philip van Acken, Evine Beursken, Linsey Roijendijk, and Pim Haselager

16.1 Introduction

BCI is an umbrella term (Nijboer et al. 2011) for several techniques where "covert mental activity is measured and used directly to control a device such as a wheelchair or a computer" (van Gerven et al. 2009). When the user performs a mental task his brain activity is measured, analyzed in real time, and used as a control signal for a device. The device then provides feedback to the user. Control is achieved through the classification of the detected activity and the mapping of this activity to an action.

In this chapter we want to explore how the insertion of a BCI in between thought and action may affect a user's sense of agency (SA), defined by Gallagher (2000) as: "The sense that I am the one who is causing or generating an action" (15). We will argue that, at least theoretically, it is possible that BCI-mediated action can leave a user uncertain as to whether or not he was the agent. We will discuss two pilot experiments we performed that illustrate how this theoretical possibility can be empirically investigated. The first experiment focused on the possibility that a user may claim to have been the agent of a BCI-mediated action, while this actually was not the case (the user had the *illusion of agency*). The second experiment examined the effect of the *transparency of the mapping* between the mental task and the performed action on this agency illusion. We will close by discussing briefly some of the potential implications of a user's uncertainty about being an agent in the process of learning to use a BCI and the potential moral and even legal implications concerning responsibility for action (Haselager 2013).

R. Vlek • L. Roijendijk • P. Haselager (✉)
Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, Nijmegen, The Netherlands
e-mail: W.Haselager@donders.ru.nl

J.-P. van Acken • E. Beursken
Department of Artificial Intelligence, Radboud University Nijmegen, Nijmegen, The Netherlands

G. Grübler and E. Hildt (eds.), *Brain-Computer Interfaces in Their Ethical, Social and Cultural Contexts*, The International Library of Ethics, Law and Technology 12,
DOI 10.1007/978-94-017-8996-7_16,

16.2 BCI and Agency: The Theoretical Issues

A person's sense of agency has been discussed increasingly over the last decade (see for example Blakemore et al. 2002; 2003; Wegner 2003a; de Vignemond and Fourneret 2004; Tsakiris et al. 2006; Pacherie 2007; Moore and Haggard 2008). On the one hand, someone might think he is doing something, though in actual fact he might merely be undergoing an event, witnessing its effects on him, but not being a cause of it. Wegner (2003a, 9) has classified this as a case of illusion of control, where someone has the feeling of doing something, but is not actually doing it. On the other hand, someone might think he is not doing something, while in fact he is. Wegner classifies this as a case of automatism: "The case of no feeling of will when there is in fact action" (Wegner et al. 2003, 9). The basic distinction is between doing something without realizing this (e.g. one's brain states being the cause of a specific event) and thinking that one is doing something without actually doing so (one's brain states not being the cause of the event).

Recently, two aspects of SA have come to be distinguished: A pre-reflective as well as a reflective SA (Gallagher 2012). Several concepts for the two different types of SA are in use and the ones best suited to give a rough, intuitive understanding are provided by Synofzik et al. (2008): pre-reflective, non-conceptual SA can be seen as *Feeling of Agency* (FoA) while a reflective, conceptual SA is labeled *Judgment of Agency* (JoA). For the FoA the most prominent models are derived from the so-called *comparator model* by Frith (1987, 2012; Gallagher 2000), while one model out of several for the JoA has been proposed by Wegner and Wheatley (1999). Although both aspects need to be studied, we decided to focus on the reflective and inferential processes involved in a participant's judgment about whether they caused or generated a BCI-mediated action.

Beginning at the end of the last century, psychologist Daniel M. Wegner published several articles on various aspects and manipulations of persons judging themselves to be the agent or 'author' of an action, i.e. judging that it was their conscious will that caused it. In their 1999 work Wegner and Wheatley looked for the sources involved in the creation of this impression. Ownership of action (or, following Synofzik et al. 2008, 'judgment of agency') arises when three preconditions are met: A thought is perceived as willed when the thought precedes the action at a proper interval (called the *priority principle*), when the thought is compatible with the reaction (*consistency principle*) and when the thought is the only apparent cause of the action (*exclusivity principle*). The *consistency principle* is especially interesting in the context of BCI. In much research, imaginary movement is used as the mental task to drive a BCI (van Gerven et al. 2009). A subject is for instance imagining left versus right hand movements. These imagined movements lead to reasonably easy to detect oscillatory neuronal patterns in right and left motor areas, which can then be used for example for the movement of a cursor on a screen. This new mapping between mental task and actuator output needs to be learned. At first participants will be uncertain which of their imagined movements will result in which exact movement of the cursor on screen. It is through double-checking the correlation between performed mental task and cursor movement that, through repetition, a consistent picture will form, reinforcing their judgment that they are the agents of the action.

It may be useful to examine how these principles help to explain the results of a BCI experiment by Lynn et al. (2010). This experiment has shown that it is possible to generate illusory intent for BCI applications – participants reported that they deliberately caused the movement of an object on a screen after being asked to try moving it as often as possible even though the movement they saw during the experiment was completely pre-rendered and allowed for no interaction. The principles help to explain this as follows. The object moved after participants allegedly began "emitting the intention of moving the line", which is in line with the *priority principle*. The object traversed the screen in a way that the participants had been led to expect through their briefing and appeared to do so consistent with their prior knowledge of, and their experience with, the BCI of Lynn et al., thus satisfying the *consistency principle*. Finally, the participants were the only visible actors, satisfying the *exclusivity principle*.

16.3 Experimental Inspiration

We wanted to illustrate how the theoretical possibility of a user's mistaken JoA could be studied experimentally. In order to do this we chose an experimental setup that remained as close as possible to one of Wegner's most vivid experiments concerning the illusion of agency. In "Vicarious agency: Experiencing control over the movements of others", Wegner et al. (2004) describe an experiment called 'helping hands'. This experiment is conducted with two people, of which one would randomly be assigned to be the participant while the other would be assigned to the role of the so-called 'hand helper'. Participants would watch themselves in a mirror while another person (the 'hand helper') stood behind them. This helper would be hidden from view in the mirror, except that he extended his hands forward on each side where normally the participant's hands would appear. The hand helper then performed a series of hand movements. Both the hand helper and the participant were wearing a headphone through which they heard instructions that sometimes were the same, and sometimes different. In Wegner et al. (2004), three different experiments were discussed. Only the first one is relevant for our purposes here. It had two conditions, a preview condition (in which the participant heard the same instructions as the hand helper) and a non-preview condition (in which the participant heard no instructions at all, but the helper still did). Results were gathered using a questionnaire with questions about how much control or conscious will the subject experienced. The answers could be rated on a Likert scale from 1 to 7. Mean vicarious control ratings were computed by taking the mean of the answers of two questions: "How much control did you feel that you had over the arms movements?" and "To what degree did you feel that you were consciously willing the arms to move?". In the preview condition, the mean vicarious control ratings reported by the subject were significantly greater (with a mean of 3.00, SD = 1.09) than in the non-preview condition (M = 2.05, SD = 1.61, t(31) = 26.8, $p < 0.02$). Although a rating of 3 on a scale from 1 to 7 (with 1 meaning no control) is still relatively low, according to Wegner the results indicate that the participants hearing the instructions just before the action "expressed an enhanced feeling that they were able to control and will the arms' movements" (Wegner et al. 2004, 841). These results show that participants in the preview condition experienced significantly more vicarious control over (indicating that the subject felt more authorship for) the movements of someone else than participants without preview. The most interesting finding of this experiment is that people experienced some sense of agency over the movements of others even though they knew someone else executed the instruction.

In this experiment two principles to make judgments about action are applied: The priority principle and the consistency principle. According to the priority principle the thought should occur before the action. In the preview condition this principle is applied (the participants thought follows the instruction), contrary to the non-preview condition in which no instruction is given and therefore no instruction-related thought occurs. The consistency principle means that the thought should be consistent with the action. This principle is applied in the consistent preview condition, but not in the non-consistent preview one (in which the instructions did not match the actions).

16.4 Experiment 1

In this experiment (see van Acken 2012 for further details), we set up a nonfunctional brain–computer interface that replaced the 'hand helper'. More specifically, we employed an electroencephalogram (EEG) hooked up to a computer that controlled the movements of a robot hand displayed on a screen in front of the subject. The signals picked up by the EEG were allegedly able to control the gestures of a virtual hand, presented on the display. In truth the participants did little more than watch a series of movies. Their EEG was measured for further research, but was not used to actually control the feedback during the experiment.

Participants would hear one of two commands, either "*thumbs up*" or "*okay*". They were told that the *vivid thought* of moving one of their hands without actually moving it (i.e. imagined movement) would be picked up by the EEG cap on their head. Upon hearing "*thumbs up*" they had to imagine moving their left hand up and down, which would in turn cause the virtual hand to perform a "*thumbs up*" gesture. Upon hearing "*okay*" from the speakers the participants had to think about moving their right hand up and down, in turn causing the virtual hand to perform the "*okay*" gesture.

Each session consisted of 60 trials, of which 30 instructed right hand imagined movement, and 30 left hand imagined movement. The order of the trials was randomized. There were two sessions in total, with two different delays between the instructions and the hand movement on screen. The session we are focusing on here had a 2.5 s delay between the onset of the audio cue and the onset of the hand movement. This delay was based on the fact that most BCIs do not react immediately but with a slight delay. Although participants were told their EEG readings would be used by the BCI to control the hand movements visible on the screen, in actual fact they were looking at fixed short movies. These movies were set to display the hand moving conforming to the audio cue 90 % of the time, and non-conforming 10 % of the time. These percentages were chosen to approximate high-level EEG-based BCI performance levels. Thus, occasionally participants would hear the "*thumbs up*" cue but see the hand perform an "*okay*" sign or vice versa. After the session participants filled in a questionnaire. We computed the participants' JoA using the following questions of Wegner et al. (2004): 'How much control did you feel that you had over the hand's movement?' and 'To what degree did you feel you were consciously willing the hand to move?'.

We performed a pilot version of this experiment with six participants (all students, three male, three female, mean age of 22.3 years, all but one right-handed, two with prior BCI experience) as part of a bachelor thesis project at the Department of Artificial Intelligence at the Radboud University Nijmegen. We examined the mean vicarious control ratings and found $M = 5.00$, $SD = 0.316$ for the 2.5 s preview condition. As the number of participants in this pilot is low ($n = 6$) we do not make any claim about significance levels. All the same, the mean vicarious control ratings reported by the participants are higher than found by Wegner et al. (2004). At the least it seems to suggest that under certain circumstances one cannot discard the option that users in a BCI context might experience an illusion of control, i.e. judge that they are the author of the act, and thereby have a sense of agency concerning actions they do not perform.

It is valuable to consider how the principles underlying judgments of agency apply to the BCI scenario of experiment 1. As in every setting the hand moved after the audio cue that served as the start signal for the participants to begin with their imagined movement, we argue that the *priority principle* holds. As the participants knew they were the only actors involved during each trial, the *exclusivity principle* holds. The answers by participants concerning how they felt about the experiment after each session seem to support this: Among the answers were phrases like "I think I did something wrong a few times, I really thought about moving my left hand but the virtual hand did the okay sign". Phrases to indicate a perceived co-authorship from the EEG system would likely have gone more along the lines of "I thought about X but *it* did Y" where *it* may be replaced by BCI, EEG, or a related term, attributing some agency to *it*. As for the *consistency principle* one might argue that it holds here since the mapping from audio cue to imagined movement to the movements of the digital hand stayed consistent except in the case of error. With all three principles accounted for – even if some might be rather weak – all sources for a JoA according to Wegner and Wheatley (1999) are present. It seems plausible to say that a BCI setting allows for stronger *exclusivity*, roughly equal or lower *priority*, and – due to the, say, novelty of the task our participants had to perform – somewhat weaker *consistency* than comparable experiments reported by Wegner et al. (2004).

16.5 Experiment 2

Experiment 2 (see Beursken 2012 for further details) focused on the effect of the transparency of the mapping between mental task and performed action on the user's JoA. The mapping between the mental task and the performed action is called transparent when the performed action conforms to the mental task. For example, normally when we think about grasping a glass (e.g. we want to grasp it) in order to grasp the glass, the thought and action 'fit' one another. In a BCI context this is rarely the case, particularly because one wants to use mental tasks that provide the largest contrast in brain signatures between them. Often left versus right hand imagined movement is used in order to for example move a cursor on the screen upwards or downwards, or having a hand-like effector opening or closing a grasp, which does not correspond directly to the imagined movement performed during the mental task.

Experiment 2 had a similar setup to experiment 1, but used two robot hands (a left and a right hand) instead of one. It had two conditions, one transparent and one non-transparent. In the transparent condition, the audio instruction, mental task and performed action conformed to each other. For instance, when the audio instruction 'left hand up' was given, the participant had to imagine moving his left hand up and down. The left virtual hand would in return move up. In the non-transparent condition, two possible audio instructions were given: '*thumbs up*' or '*okay*'. When the instruction 'thumbs up' was given, the participant had to imagine moving his left hand up and down. The virtual hands would in return make a thumbs up sign. When the instruction 'okay' was given, the participant had to imagine moving his right hand up and down. The virtual hands would in return make an okay sign. As in experiment 1, the gestures of the virtual robot hands were preprogrammed, which means that the participants were not in control, even though they might think they were.

We performed a pilot experiment to investigate transparency as part of a bachelor thesis project at the Department of Artificial Intelligence at the Radboud University Nijmegen. There were eight participants (four females and four males), between the ages of 19 and 25. None of them reported to have experience with BCI (though three reported to have experience with EEG). Six of them were right-handed, and two were left-handed. The experiment had a within-subject design in which each subject performed a session of each condition (the order was randomized and counterbalanced). Each session contained 60 trials (30 left hand and 30 right hand, randomized over the experiment). As in experiment 1 we also chose to simulate a BCI performance of 90 % correct.

As in experiment 1, the user's JoA is computed using the following questions: 'How much control did you feel that you had over the hand's movement?' and 'To what degree did you feel you were consciously willing the hand to move?'. We compared these ratings (one for each condition) using a paired-samples t-test. Mean vicarious agency ratings were higher in the transparent condition ($M = 4.188$, $SD = 1.710$) than in the non-transparent condition ($M = 3.688$, $SD = 1.557$), but this difference was not significant ($t(8) = 15.28$, $p = 0.17$ (2-tailed)). Thus, the transparency of the mapping between the mental task and the performed action did not significantly influence the user's JoA. However, the ratings are found to be in the right direction (higher in the transparent condition than in the non-transparent condition). The results support those of experiment 1, in that they are higher than those reported in Wegner et al. Again we refrain from making any claims about significance here because of the low number of participants.

16.6 Differences Between the Two Experiments Potentially Relevant to JoA

In order to obtain a better understanding of how the context of a BCI setting can be of influence on a user's JoA, it is interesting to consider why the vicarious agency ratings in experiment 1 have a higher mean ($M = 5.00$) than in the transparent and non-transparent condition of experiment 2 ($M = 4.19$ and $M = 3.69$, respectively). This difference is especially interesting since experiment 1 used a non-transparent set-up. To explain these results we noticed the following differences that might have played a role:

Showing live recordings of the data recorded by the electrodes to the participant. In experiment 1, the live recordings of the electrodes were shown to the participant during the setting up of the EEG equipment. The participants were asked to blink their eyes and clench their teeth, while the EEG waves were shown on the screen in front of them. In this way they saw direct feedback from their brain. This may have had a significant influence on the participants. Even though this does not immediately show that the BCI performs well, it might help to convince the participant that the EEG signals are used for controlling the BCI output. In experiment 2, no such live recordings were shown to the participants and this might result in lower reinforcement of the suggestion that the user actually controls the output.

Availability of the instructions during the experiment. In experiment 1, the instructions were available on paper during the experiment, allowing the participant to look at the paper to see what needed to be imagined when a certain instruction was heard. Showing the instructions does not make the mapping between the mental task and performed action more transparent, because the mental task and performed action still do not conform to each other. However, showing the instructions may compensate for the lack of transparency between the mental task and performed action by reducing the memory load.

Number of virtual robot hands. In experiment 2 two virtual robot hands were used (to associate with left and right hand imagined movement), while in experiment 1 only the right hand was used. One might consider that one versus two virtual robot hands had an influence on the JoA and that two virtual robot hands can better reflect reality, since participants have two hands and both hands were used for imagination. This would suggest that the participant might have felt more control (because the consistency principle can be better applied here) and would therefore have given higher ratings. However, the results contradict this; in experiment 1 higher ratings were found than in experiment 2. We are at a loss how to explain this except for the suggestion that the difference in results is caused by one of the other differences between the experiments.

To summarize, probably the first difference (showing live recordings) has a large impact on the JoA. The difference in the number of virtual robot hands gives slightly contradictory results, since we expected the participant to feel more control over two virtual robot hands than over one. Showing the instructions during the experiment (difference number two) could have helped the participant to internalize the mapping between the mental task and the performed action better. This could give the participant the idea of having more control. Unfortunately more than one difference is found between the experiments, so that we cannot attribute the difference in results to one specific cause. We analyzed these issues in such detail to provide an example of how subtle details of the context in which a BCI is used may be of influence on a user's JoA. Further studies need to take this into account.

16.7 Conclusion

First of all, our pilot studies seem to indicate that it is possible to measure a user's JoA in a BCI context. Furthermore, in both experiments a JoA rating was found that was relatively high (i.e. above the 50 % point of the 1–7 Likert scale). While it thus seems possible to evoke an illusion of agency, two important questions remain: (1) What are the potential implications of unjustified JoAs? And (2) could a manipulation of a user's JoA be beneficial to BCI use and, if so, under what circumstances?

Imagine a user – let's call him Frank – trying to perform an action through the assistance of a BCI. Frank is imagining his left hand moving, in order to have a robot hand picking up a cup of hot coffee so that a person (Louis) sitting nearby can get it. Something goes wrong and the coffee spills over Louis, resulting in damaged clothes and slight skin burns in the process. Does Frank feel responsibility – as distinct from being (legally) responsible – for the outcome of his attempt? If our investigations are on the right track, it may be that Frank may feel that he himself did something wrong (e.g. in his performance of the mental task), while actually he (his brain states) played no causal role in the unfortunate outcome. That is, it may have been that he performed the mental task correctly, but some aspect of the BCI was not working properly. Though Frank did not really do it, he may feel he did, and feels the guilt and perhaps pays damages accordingly. Importantly, there may be inconspicuous aspects of the context in which Frank is using the BCI that may strongly influence his feeling of responsibility. Our analysis and experiments tentatively suggest that BCI users may be well advised to carefully (re-)consider whether their first assessment of responsibility is correct.

From a legal perspective, there might be important implications as well (see Tamburrini 2009 for a brief discussion of potential liability issues). According to the legal dictum used in American criminal law 'actus non facit reum nisi mens sit rea', the act does not make a person guilty unless the mind is also guilty. An important condition for Frank having a guilty mind ('mens rea') is that he is aware that he is acting in a specific way, while doing it. An act performed by sleepwalking, for instance, is considered to be involuntary, which may exempt a person from at least part of his responsibility for the outcome. In the case of acts leading to undesired consequences (such as in the case of Frank with the spilling of coffee), a legally important criterion is negligence. Negligence arises when the accused unintentionally committed the criminal act without exercising the care that a reasonably prudent person would exercise in similar circumstances. As some of the participants indicated after the experiments we performed, Frank might be blaming himself for not exercising reasonable care while using the BCI (e.g. reports himself to have been distracted), thereby admitting to a legally relevant error that he actually did not make.

Regarding the second issue, perhaps surprisingly there could actually be some benefits of a higher JoA in BCI users. For instance, an unjustified JoA could have stimulating effects in the often difficult initial phase of learning to use a BCI. It might help participants to pass the stage where one has no idea what one is doing or whether what one is doing is right. Moreover, a heightened JoA might evoke some sense of ownership towards the BCI, which could trigger the so-called *mere ownership effect* (Beggan 1992). That effect states that the attitude towards an object gets more positive if one owns the object compared to the attitude towards the same object if one does not own it. As such *the mere ownership effect* could in turn lead to a more positive evaluation of the BCI by the user. A user that likes his BCI might be more forgiving towards errors. In addition, positive affect (alongside glucose and resting) is known to replenish mental resources to an extent and motivate subsequent task engagement (e.g., Tice et al. 2007; Thoman et al. 2011); training or using a BCI would feel less tiresome for users who view *their* BCI in a positive light. For both the positive and negative reasons, and for the theoretical interest of sense of agency in its own right, we would like to suggest that it is important to further study the effects of BCI use on judgments of agency.

References

Beggan, J.K. 1992. On the social nature of nonsocial perception: The mere ownership effect. *Journal of Personality and Social Psychology* 62(2): 229–237.

Beursken, E. 2012. *Transparency in BCI: The effect of the mapping between an imagined movement and the resulting action on a user's sense of agency*. Bachelor thesis, Department of Artificial Intelligence. Radboud University, Nijmegen, The Netherlands. http://www.ru.nl/artificialintelligence/bachelor/voltooide-bachelor/

Blakemore, S.-J., D.M. Wolpert, and C. Frith. 2002. Abnormalities in the awareness of action. *Trends in Cognitive Sciences* 6: 237–242.

Blakemore, S.-J., D.A. Oakley, and C.D. Frith. 2003. Delusions of alien control in the normal brain. *Neuropsychologia* 41(8): 1058–1067.
de Vignemond, F., and P. Fourneret. 2004. The sense of agency: A philosophical and empirical review of the 'who' system. *Consciousness and Cognition* 13: 1–19.
Frith, C.D. 1987. The positive and negative symptoms of schizophrenia reflect impairments in the perception and initiation of action. *Psychological Medicine* 17(3): 631–648.
Frith, C. 2012. Explaining delusions of control: The comparator model 20 years on. *Consciousness and Cognition* 21: 52–54. doi:10.1016/j.concog.2011.06.010.
Gallagher, S. 2000. Philosophical conceptions of the self: Implications for cognitive science. *Trends in Cognitive Sciences* 4(1): 14–21.
Gallagher, S. 2012. Multiple aspects in the sense of agency. *New Ideas in Psychology* 30(1): 15–31. doi:10.1016/j.newideapsych.2010.03.003.
Haselager, W.F.G. 2013. Did I do that? Brain-computer interfacing and the sense of agency. *Minds and Machines* 23(3): 405–418. doi:10.1007/s11023-012-9298-7.
Lynn, M.T., C.C. Berger, T.A. Riddle, and E. Morsella. 2010. Mind control? Creating illusory intentions through a phony brain-computer interface. *Consciousness and Cognition* 19: 1007–1012. doi:10.1016/j.concog.2010.05.007.
Moore, J., and P. Haggard. 2008. Awareness of action: Inference and prediction. *Consciousness and Cognition* 17: 136–144.
Nijboer, F., J. Clausen, B.Z. Allison, and P. Haselager. 2011. The asilomar survey: Stakeholders' opinions on ethical issues related to brain-computer interfacing. *Neuroethics*. doi:10.1007/s12152-011-9132-6.
Pacherie, E. 2007. The sense of control and the sense of agency. *Psyche* 13(1): 1–30.
Synofzik, M., G. Vosgerau, and A. Newen. 2008. Beyond the comparator model: A multifactorial two-step account of agency. *Consciousness and Cognition* 17(1): 219–239. doi:10.1016/j.concog.2007.03.010.
Tamburrini, G. 2009. Brain to computer communication: Ethical perspectives on interaction models. *Neuroethics* 2(3): 137–149. doi:10.1007/s12152-009-9040-1.
Thoman, D.B., J.L. Smith, and P.J. Silvia. 2011. The resource replenishment function of interest. *Social Psychological and Personality Science* 2(6): 592–599. doi:10.1177/1948550611402521.
Tice, D.M., R.F. Baumeister, D. Shmueli, and M. Muraven. 2007. Restoring the self: Positive affect helps improve self-regulation following ego depletion. *Journal of Experimental Social Psychology* 43(3): 379–384. doi:10.1016/j.jesp.2006.05.007.
Tsakiris, M., G. Prabhu, and P. Haggard. 2006. Having a body versus moving your body: How agency structures body-ownership. *Consciousness and Cognition* 15: 423–432.
van Acken, J.-P. 2012. *Tracking the sense of agency in BCI application*. Bachelor thesis, Department of Artificial Intelligence. Radboud University, Nijmegen, The Netherlands. http://www.ru.nl/artificialintelligence/bachelor/voltooide-bachelor/
van Gerven, M., J. Farquhar, R. Schaefer, R. Vlek, J. Geuze, A. Nijholt, N. Ramsey, et al. 2009. The brain-computer interface cycle. *Journal of Neural Engineering* 6(4): 041001. doi:10.1088/1741-2560/6/4/041001.
Wegner, D.M. 2003a. *The illusion of conscious will*. Cambridge: MIT Press.
Wegner, D.M., and T. Wheatley. 1999. Apparent mental causation: Sources of the experience of will. *American Psychologist* 54: 480–492. doi:10.1037/0003-066X.54.7.480.
Wegner, D.M., V.A. Fuller, and B. Sparrow. 2003. Clever hands: Uncontrolled intelligence in facilitated communication. *Journal of Personality and Social Psychology* 85: 5–19. doi:10.1037/0022-3514.85.1.5.
Wegner, D.M., B. Sparrow, and L. Winerman. 2004. Vicarious agency: Experiencing control over the movements of others. *Journal of Personality and Social Psychology* 86: 838–848. doi:10.1037/0022-3514.86.6.838.