

have the illusion of conscious will so systematically and so pervasively? Perceptual illusions are explicable as unusual violations of the conditions under which our sensory processing are veridical, but attributions of free will are scarcely unusual, and an explanation is required. It is hard to resist attributing autonomy to others, even when we see the mechanics of reason come apart before our eyes. Anyone who has had day-to-day encounters with someone suffering from obsessive/compulsive disorder will have had the impulse to blame the sufferer for irrational actions committed in the course of their otherwise normal conduct and discourse. If we have no Cartesian freedom of the will, why do we have so fierce an inclination to attribute autonomy to ourselves and others? What function, what cognitive causal role, do such beliefs have that might help to explain their emergence and retention in the human psyche, and why do we have them *consciously*? Wegner offers an answer to the first of these twinned questions. I will offer another.

Wegner sketches this answer: Our conscious illusions of autonomous action inform us about ourselves and prompt feelings of moral responsibility and guilt, which influence our subsequent actions. That answer seems correct so far as it goes, but inadequate to the question. One could conceivably be perfectly aware of one's own actions without having the sense that one does them autonomously. Wegner's proposal does not explain why we attribute *others'* actions to their autonomous intentions with nearly the same force and immediacy of our self-attributions; nor does it explain why knowledge of action need be conscious – but neither will I.

Rather, here is another conjecture: *The implicit assumption of freedom of the will is essential to learning. If we did not at least unconsciously assume our own actions to be autonomous, we could not learn the effects of our own actions; and if we did not assume the same of others, we could not learn the effects of our own actions by observing theirs. If, in action taken or observed, the application of that assumption is conscious, we must have the illusion of conscious will.*

Consider scientific inference from observational, non-experimental, data. There are several possible explanations for a correlation observed among two kinds of events for which instances of one kind precede those of the other: Events of the first kind may cause the second; or some third factor or factors may influence both kinds of events; and there are still other possibilities. For concreteness, consider an association between smoking and lung disease, which could be explained by at least two different causal structures:

1. *Smoking* → *Lung Disease*
2. *Smoking* ← *Unknown* → *Lung Disease*

To make a reasonable causal inference, one must have grounds to exclude the second explanation. One rarely does, and that is why observational science is hard. Experimentation tends to eliminate alternative explanations of data. What makes an experiment an experiment is that acting from outside the system under study, the experimenter determines the value of the causal variable, or determines its probability distribution. If the experimenter fixes or randomizes the value of the causal variable in each case, and does so by a method not influenced by other features not under the experimenter's control, then there is no confounding. If we force someone – or an entire population – not to smoke, then we eliminate confounding, and, if smoking does not cause lung disease then these two variables are uncorrelated in the experimental results. (For mathematical details, see Pearl 2000; Spirtes et al. 2001; and for a philosophical exposition, see Woodward 2003.)

Independent manipulation does not make causal learning possible, but it makes it enormously easier to make accurate causal inferences. Whatever the circumstances, if one does not impose the premise – warranted or not – that the association of putative cause and effect is not produced by other common causes of both, the inference to causation is wanton.

For our inner workings – the unconscious, biological algorithms of thought – to allow that actions have unknown causes would be precisely for them to allow that those unknowns might also cause the immediate and slightly more remote events that we take to be

effects of actions; action and event would be potentially confounded and no causal inference would be possible in everyday life, just as no causal conclusions are possible in ill-designed, confounded, scientific experiments or in poorly designed observational studies. So, unconsciously at least, to be intelligent in the way we are, we *must* presuppose autonomous actions – and to make correct causal inferences, actions and their effects must for the most part actually be unconfounded by common causes. An organism that did not so assume might learn by association, but its ability to plan and foresee the effects of interventions in the world would be severely limited. Daniel Povinelli (2000) and Tomasello and Call (1997) give evidence that our nearest biological neighbors are limited in these respects, while Gopnik et al. (2004) give evidence that even quite young children make comparatively sophisticated causal inferences from data in which passive correlations and effects of interventions are combined. If, from whatever causes, the assumptions of our inner processes that lead to action are consciously manifested in the very instance of action or in the perception of action in others, we will have the conscious sense of autonomous agency, of freedom of the will. And we do. We think immediately that our actions cause the observed effects, and nothing else causes both our actions and the observed effects. Usually, we assume the same of others, and if we did not then we could not learn causal relations from their actions and the events that follow them.

ACKNOWLEDGMENTS

Thanks to Alison Gopnik for helpful discussions, to Sam Buss for the best argument, and to the Office of Naval Research for support of research.

The elusive illusion of sensation

Valerie Gray Hardcastle

Department of Science and Technology in Society, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061-0247. valerie@vt.edu
<http://www.mind.vt.edu/>

Abstract: The sensation of will is not the same thing as the will itself any more than the sensation of hunger is the same thing as being devoid of nutrients. This is not a really surprising claim, but it is the only claim to which Wegner is entitled in his book.

When I feel hunger pangs, am I feeling genuine hunger, or am I feeling “merely” the sensation that accompanies real hunger, a purely physiological state? If the latter, then hunger pangs must be some sort of illusion, a stand-in for states we cannot access consciously. When our bodies infer that they need more nutrients, we feel hungry. However, as the popular press makes very clear, we are often wrong about this inference and consequently feel hungry when we aren't really.

This meditation on hunger parallels what Wegner (2002) says about our sensations of willing an action. The sensation of willing isn't actually doing anything; it certainly isn't causing our bodies to behave in any particular way. Instead, the sensation is “merely” telling us that (we think) our own psychological states are driving our bodies.

Is this conclusion so surprising? I grant that we generally talk and think about the will in very sloppy terms, but when we get right down to it, do we really believe that the sensation of willing just is the will itself? I submit that we do not; we believe, if we have ever even thought about these matters before, that the sensation informs us about the sort of actions we are performing. If we feel the force of our will, then we believe that we, in some important and fundamental sense, are the causal agents responsible for what we are doing. The sensation of will isn't the will itself any more than the sensation of hunger is the same thing as being devoid of nutrients, or the sensation of warmth is heat itself, or the smell of a rose is the rose itself. In each case, our sensations tell us something about the world out there (or in here); they indicate or represent to us the way the world is (or we take it to be).

Wegner provides us with case after case of how our sensations of will are mistaken, how we sometimes do things ourselves but attribute these actions to others, how we sometimes think we are doing things ourselves, but we aren't. He is right; our sense of will is sometimes – maybe a lot of the time – misleading.

But so what? What, if anything, does this tell us about freedom of the will – the actual will, not what we sense as a marker for the will? Not much. In order to know something about the actual will, for example, whether it exists in any interesting sense, we would have to know how the sensation of will connects up with either our underlying psychology or our underlying physiology or both. However, unlike the case of hunger, in which we know a lot about the connection between various levels of hormones in our blood stream and wanting to eat, we know very little about what the sensation of will actually reflects. Maybe it does mirror a genuine self in the brain: our central control that initiates or at least approves our purposeful behaviors. Maybe it doesn't. But knowing that our conscious sensations of will are sometimes mistaken doesn't shed any light on this topic.

We know some actions happen to us – I sneezed in the middle of lecture – and others have a psychological reason behind them – I raised my hand in the middle of lecture. We can tell the difference between these sorts of activities, both from the inside, as it were, and from the outside. But what is this difference? Is it just that the latter is accompanied by a sensation of will and the former isn't? Is it just that we explain the latter in terms of beliefs and desires and the former in terms of physiology? Or does the latter occur as the endpoint in a causal chain mediated by my own psychological states, whereas the former doesn't? I think that no matter what one's metaphysical stripe, one would have to agree with the last suggestion: What differentiates willed actions from actions that are not willed is the causal history of the action. Willed actions flow from or through my psychological states in ways that unwilled actions don't.

But if this is the case, then in what sense is our sensation of conscious will an illusion? Our sensation serves to differentiate which actions flowed from or through our psychological economy from those that did not. It may get it wrong once in a while; it may get it wrong lots of times. Nevertheless, the sensation is reflecting something real, as real as our bodies' need for nutrients. The important question is what exactly is that sensation reflecting.

Wegner wants to argue that we don't really have selves, that our sensation of selfhood, too, is just another inference our bodies and brains make about what we are doing in order to explain our selves to ourselves. He wants to argue that we have this whole edifice of illusory constructions about our own psychologies from which the sensation of will flows. He wants to argue these things, but he can't. He can't because he doesn't get below the sensations to learn what is really going on. He has "Just So" stories about how selves might work, but so do a lot of people. Unless and until we can connect our sensations to actual physiological or deeper psychological workings, it will be hard to claim that our sensation of will is illusory in any interesting sense.

The sense of conscious will

Gene M. Heyman

Behavioral Psychopharmacology Research Laboratory and Psychiatry, McLean Hospital and Harvard Medical School, Belmont, MA 02478.
gmheyman@mclean.harvard.edu

Abstract: Wegner's conclusion that conscious will is an illusion follows from a key omission in his analysis. Although he describes conscious will as an experience, akin to one of the senses, he omits its objective correlate. The degree to which behavior can be influenced by its consequences (voluntariness) provides an objective correlate for conscious will. With conscious will anchored to voluntariness, the illusion disappears.

When an object, say a boat on the water, moves away, its retinal image decreases in size. However, instead of experiencing the boat

as shrinking, the viewer experiences it as receding into the distance. This could be called an illusion; the retinal image is getting smaller, not further away. However, to say that this is an "illusion" is to ignore the determinants of object constancy. When the viewer's understanding of boats and the three-dimensional world are included in the analysis of visual experience, the correlation between a shrinking retinal image and the perception of a constant sized but increasingly distant object is perfectly understandable. Or, to put it another way, to say that one of the visual constancies is an "illusion" is to overlook that there is more to vision than the retina.

Wegner's treatment of conscious will (Wegner 2002) is rather like trying to account for object constancy while limiting the analysis to the retina. He emphasizes that conscious will is an inference and that its contents often do not match up well with the actual factors that cause voluntary action. For instance, we may be aware of the intention to raise our hand (or assume this intention after the fact), but not be aware of the determinants of this intention or of having made an inference. From these "discrepancies," Wegner concludes that conscious will is an illusion. However, as in the object constancy example, a more complete account of the input eliminates the illusion.

Object constancy is about the fact that we live in a three-dimensional world and that when objects move, they usually do not change shape. What is conscious will about? What is its stimulus? The answer cannot be found in a textbook (as with the perceptual constancies), but it is familiar and easily identified.

As documented by Wegner, conscious will's domain is behavior, in particular our own behavior. Just as perception tracks dimensions of the external world, conscious will tracks the important fact that our own activities vary in the degree to which they are influenced by consequences (e.g., rewards, incentives, punishments, and the like), by the values we adopt, and by new information. Some activities are immune to these factors, whereas others are easily modified by just a hint of praise or disapproval. For instance, consider the different causal relations relating to a patellar reflex and learning to kick a ball, blushing and the decision to wear rouge, a defensive blink and a conspiratorial wink at a friend. The second activity in each comparison we call voluntary, and the first we call involuntary. The distinction is not a matter of free will versus determinism. Antecedents govern voluntary and involuntary acts. Rather, the mediating neural architecture and nature of the antecedents differ. Differences in neural connections allow for variation in the degree to which activities are influenced by experience and the contents of consciousness. The distinction also does not depend on intentions or other subjective reports. We can be conscious of involuntary acts (I know I am going to blink, but I can't help it), and as Wegner's literature review ably demonstrates, we can be unaware of voluntary acts. In other words, voluntariness (susceptibility to consequences) provides an objective basis for subjective experience, just as the conservation of an object's shape and size while moving provides a basis for perceptual constancy.

Wegner acknowledges that behavior varies with regard to its susceptibility to consequences (e.g., the ear wiggling discussion, Wegner 2002, pp. 31–34), and also acknowledges that voluntary actions are the usual focus of conscious will. However, these observations are made in passing, and his analysis proceeds without any further discussion of the objective basis for the sensation of "doing something." Given this omission it is understandable that he concludes that it is an illusion. This is not to say that conscious will is a literal reflection of susceptibility to rewards. For instance as Wegner's discussion of automatic processes (2002, pp. 56–59) demonstrates, many learned, reinforced actions can move out of awareness.

Leaving out the objective correlates for conscious will leads to empirical and logical problems. An empirical shortcoming is the de-emphasis of the contribution that conscious will makes to voluntary action. Often Wegner seems to be saying that conscious will is no more than an after-the-fact frill, at best useful for a kind of moral bookkeeping (see below). I am not sure that this is what he