processes that cause the behaviors. At least one line of research (that on the interpretation of experimental studies of mental imagery, e.g., Pylyshyn 1981) has provided evidence that the content and dynamics of people's consciously experienced imagery is a result of, rather than a cause of, what people know about the situation being imagined (as well as the way that the task of imagining is interpreted).

Studies of mental imagery provide a clear example of the illusory nature of explanations based on the conscious experience of one's cognitive process (for a detailed discussion of this issue, see Pylyshyn 2003b). One's experience of the *form* of images and of why they unfold the way they do provides one of the most misleading sources of explanations of mental processes. The experience of "seeing" events unfold in one's "mind's eye," and thereby of seeing *why* some operations are more difficult or take longer than others, is so compelling that it is almost impossible to discount. The fact that the conscious experience of visual imagery is similar to that of perception (presumably because of the involvement of some of the same brain mechanisms) suggests to many people that the representations involved in imagery must themselves *resemble* the content of the experience (viz., that they consist of picture-like displays) or that images are constrained in their dynamics by principles similar to those that govern the world being imagined (e.g., that they "rotate" while rigidly retaining their shape). Yet the inference from the form of the experience of imagining to the picture-theory (or a theory that claims we have a dynamic model of the world in our head) is based on unsupported assumptions, such as that a brain state responsible for the conscious experience of seeing must itself resemble what is seen, or that the brain is so constituted that images are required to follow principles similar to those that govern processes in the physical world.

As in the case of experienced volition, our conscious experience appears to provide a natural explanation of *why* certain behaviors occur. In mental imagery we not only have the experience of "seeing" but we also have the experience that certain patterns of the ensuing behavior are *caused by* properties of the representations that we consciously experience. For example, it seems clear why it takes us longer to report details in a "small" image than in a "large" one; our experience tells us that this is because the details are "harder to see" when the image is smaller. Similarly, it seems obvious why it takes us longer to switch our attention between two imagined objects when they are imagined as being further apart; our conscious experience shows that this is because attention takes longer to move a greater distance across the surface of the image. Likewise, it is no puzzle why we find it more difficult to see the outer edges of our mental image; our experience shows us that this is because our "mind's eye" has a certain visual angle and when things get near the periphery they are harder to discriminate, just as things are harder to see in the periphery of vision. It thus seems that many properties of mental imagery, including why certain results are obtained in imagery experiments, can be explained by simply attending to the experience and seeing for yourself how the process happens.

However, the explanations suggested by conscious experience can easily be shown to be specious in examples such as the ones cited above. Even though our causal mental process may go through a sequence that corresponds to the sequence that we experience, it does not in general proceed that way for *the reasons suggested by the conscious experience* (for more on this, see Pylyshyn 2002; 2003a). The way our imagery unfolds – the sequence it goes through when we imagine certain events – is consistent with objective measures such as reaction times, but it cannot explain them. The experience of taking longer to scan greater imagined distances does not explain the reaction time observations, because the principle that it takes longer to travel a greater distance applies only to real motion over real physical distances, not to phenomenological motion, which can follow any principle one wishes (try imagining that your attention hops from place to place in your image without taking time that increases with distance). The real reason that our imagery goes through the sequence it does is, in many cases, simply that we *make* it go through that sequence because that is the sequence we expect in the situation we are imagining. To imagine something means to recreate what one believes would happen in the situation one is imagining. In other words, what we experience as arising from properties of the image itself is actually a consequence of our knowledge of how things would work in the imagined world. Evidence for this is that if we change what people believe would happen in the imagined situation, the observations also change predictably (see Pylyshyn 1981 for examples). What explains the behavior in these cases is not some principle that governs the dynamics of our image, as suggested by the conscious experience of watching the imagery unfold autonomously in one's mind's eye, but rather our (generally tacit) knowledge of the situation we are imagining (together with the psychophysical ability to simulate the sequence). It is in this sense that the conscious experience of mental imagery might be viewed as "illusory," though a better way to characterize it is that the experience of mental imagery provides a misleading explanation of why certain patterns of behavior occur.

A social psychologist illuminates cognition

Amir Raz[a] and Kim L. Norman[b]

[a]*Columbia University College of Physicians and Surgeons and New York State Psychiatric Institute, New York, NY 10032;* [b]*Department of Psychology, Barnard College, New York, NY 10027.* **ar2241@columbia.edu kn2010@barnard.edu**

Abstract: Sprinkled with humor, social psychology illuminates cognition in Wegner's beautifully written and cleverly crafted book. However, scantily exploiting such themes as psychopathology, development, and neural correlates of consciousness, Wegner's account does not fully project into cognitive neuroscience. Broaching the topic of self-regulation, we outline neurocognitive data supplementing the notion that voluntariness is perhaps more post hoc ascriptions than bona fide introspection.

Combining phenomenology with empirical data, Wegner, a social psychologist, skillfully elucidates the relationship between willed action and its underlying representations, taking the reader from the labyrinths of parlor magic into the realms of hypnosis. It is possible to gain insights into both healthy and pathological function by examining the healthy individual under atypical conditions. Social psychologists have regularly and successfully employed this research model, recruiting such tools as suggestion and deception into their research arsenal. Whereas researchers in social psychology may "push" normal individuals towards the pathological spectrum in their efforts to illuminate behavior, cognitive neuroscientists have largely subscribed to the opposite approach (i.e., studying patients with specific brain lesions, trying to understand the nonpathological or healthy brain). For example, that more attention should be given to the investigation of healthy individuals driven towards the neuropsychological domain is evident in the recent contributions of social psychology to cognitive science (Wegner 2003a) and the impact of transcranial magnetic stimulation (TMS) (George 2003). However, Wegner's (2002) account only scantily touches on psychopathology and does not fully exploit the role of development or recent knowledge concerning the neural correlates of consciousness. As a case in point, his chapter (chapter 8) on hypnosis warrants a closer look.

Hypnosis can undoubtedly bring about an observed alteration in volitional control over behavior and offers much insight into substrates of authorship. Whereas historically it has been assumed that these behaviors were indeed unintentional and that hypnosis occurred when a subject surrendered control to the hypnotic operator (Woody & Bowers 1994), in later years an alternative view emerged proposing that although hypnosis may cause a subject to be unaware of having a particular intention, these responses are in fact intentional (Kihlstrom 1992c). These two opposing views

of the volitional status of suggested behavior have not only become the subject of passionate debate but constitute the crux differentiating theories of hypnosis (Kirsch & Lynn 1998a). Hypnosis data drawing on cognitive science and neuroimaging have provided significant insights into this conundrum.

It is not a coincidence that some practitioners prefer the term "self-hypnosis" to "hypnosis" (cf. Olness & Kohen 1996). Participants in hypnosis studies generally wish to be hypnotized and therefore consent to fill the hypnotic role and follow suggestions. Their compliance differs from that of a voluntary response to a request in that they must make plans not only to execute a suggested movement, but also to concurrently interpret the movement as non-volitional. Indeed, there are data supporting this mental process (Silva & Kirsch 1992). However, whether or not hypnotic responses are intentional, it is important to remember that they are experienced as involuntary by the subject. As it is likely that these responses are a product of both intentional and automatic elements, the issue becomes more a question of whether the response is elicited *intentionally* or *attentionally* (e.g., Raz & Shapiro 2002).

There are data showing that highly hypnotizable individuals can eliminate involuntary and ballistic effects (e.g., Stroop interference) following a specific posthypnotic suggestion (MacLeod & Sheehan 2003; Raz et al. 2002; 2003b; Schatzman 1980). When they do, specific brain changes related to this effect occur (Raz 2004). Furthermore, there are now genetic findings concerning individual differences that might relate to the distinction between highly and less hypnotizable people (Raz et al. 2003a; 2004; in press) as well as evidence that hypnotic inductions might lead to "behavioral lesions" reminiscent of actions following veridical lesions (e.g., stroke) (Raz 2004). Indeed, the heritability of hypnotizability is among the highest of any psychological individual-difference measure identified to date (Morgan 1973; Morgan et al. 1970) and neuroimaging findings associated with such hypnotic and attentional modulations consistently implicate differential activation patterns in the anterior cingulate cortex (ACC) (Fan et al. 2003; Raz et al. 2003a; 2004; in press).

A popular theory of cognitive control proposes that the ACC is part of a network involved in handling conflict between neural areas. While some researchers view the ACC through the lens of a conflict-monitoring model (Botvinick et al. 2001; Cohen et al. 2000), others construe it as a regulation model engulfing broader processes of consciousness and self-regulation, including executive attention and mentation (Bush et al. 2000). Consistent with the importance of the ACC to normal self-monitoring, there are syndromes of abnormal agency that occur with extensive lesions of the ACC and associated midline frontal cortex whereby a patient interprets the actions as caused by an outside force (Goldberg 1985). The ACC is well-situated to mediate between limbic motivational influences and the adjacent supplementary motor areas, and lesions associated with ACC and medial frontal regions have been documented to produce akinetic syndromes, in which patients do not engage in actions despite being quite capable of doing so (Damasio & Van Hoesen 1983). With their ACC impaired, these patients appear to lack motivation to act. Towards this end, psychosurgery sometimes aims for the ACC to alleviate chronic pain or decrease the symptoms of anxiety, as such interventions typically decrease the patient's concern over life problems (Rainville et al. 1997).

The illusion of conscious will can be also harnessed towards a low-cost and noninvasive therapeutic means. For example, hypnotic interventions have been used to alleviate tic symptoms in individuals diagnosed with Tourette syndrome (TS) (Crawford 1992; Culbertson 1989; Kohen 1995; Kohen & Botts 1987; Lindner & Stevens 1967; Young & Montano 1988; Zahm 1987). Hypnotic suggestion is believed to engage self-regulatory mechanisms (Ray & Tucker 2003), and, whereas effortful control can evanescently suppress TS symptomatology, rendering self-regulation a lens by which to view TS formulation, the fact that volitional as well as involuntary control of behavior can be interrupted and modified by external suggestion proposes that, at least under appropriate conditions, hypnotic influence may engage mechanisms of control at an elementary level. By understanding the substrates of these processes, therefore, we may better understand not only the interesting phenomenon of conscious will, but mechanisms of self-regulation. This is particularly appealing in the context of human development, wherein studies have shown that the sense of control over actions becomes stronger with age. In this regard, studies of hypnotic susceptibility have repeatedly shown that children are more hypnotizable than adults (London 1965; Olness & Kohen 1996) and more readily attribute the cause of their actions to an external source, suggesting that the separation of action from authorship is perhaps more potent in younger age. The maturation of self-regulatory mechanisms across development is instructive in this sense, because prefrontal brain development reflects changes in perception of control over action as well as thought and emotion and may lead to a more complete understanding of the correlates of conscious will (Bronson 2000).

In conclusion, Wegner's book is a delightful composition and a fine demonstration of how cognitive science can learn from the insights of an accomplished social psychologist. Although we would have liked to see a more rigorous treatment of relevant psychopathology and, particularly, data concerning the neural correlates of consciousness, books take time to prepare and some of the data we cite here were probably unavailable as Wegner was putting pen to paper. Apropos, Christof Koch's latest, *Quest for Consciousness* (2004) nicely complements Wegner's efforts on these points.

Conscious will in the absence of ghosts, hypnotists, and other people

Johannes Schultz[a], Natalie Sebanz[b], and Chris Frith[c]

[a,c]*Wellcome Department of Imaging Neuroscience, University College London, London, WC1N 3BG, United Kingdom;* [b]*Max-Planck Institute for Psychological Research, 80799 Munich, Germany.*

j.schultz@fil.ion.ucl.ac.uk sebanz@psy.mpg.de

c.frith@fil.ion.ucl.ac.uk http://www.fil.ion.ucl.ac.uk

http://www.mpipf-muenchen.mpg.de/cgi-bin/mitarbeiter.cgi?name=SENA&filetype=personal&language=g

http://www.fil.ion.ucl.ac.uk/princdir/frith.html

Abstract: We suggest that certain experiences reported by patients with schizophrenia show that priority, consistency, and exclusivity are not sufficient for the experience of willing an action. Furthermore, we argue that even if priority, consistency, and exclusivity cause the experience of being the author of an action, this does not mean that conscious will is an illusion.

Wegner (2002) discusses an impressive variety of phenomena demonstrating that when the three conditions, priority, consistency, and exclusivity are met, an action feels willed, whereas when one or more do not apply, the cause of an action is attributed to forces other than the self. He convincingly shows that the feeling of conscious will can be erroneous, such that a person can either believe he was the author of an action even though he was not, or that he can believe he was not the author while in actual fact he was. The strongest version of Wegner's claim would be that priority, consistency, and exclusivity are both necessary and sufficient for the experience of willing an action. However, we suggest that certain experiences reported by patients with schizophrenia show that priority, consistency, and exclusivity are not sufficient for the experience of willing an action.

Patients with delusions of control report that their actions, even quite trivial actions, are being controlled, not by themselves, but by some alien force. Patients report such abnormal experiences even though they have the prior intention to make the action, the action made is consistent with their intention, and there is no obvious ambiguity about who is making the action. We have suggested elsewhere (Hohwy & Frith 2004) that what is missing is an aspect of