

9 The Hazards of Claiming to Have Solved the Hard Problem of Free Will

Azim F. Shariff
Jonathan Schooler
Kathleen D. Vohs

She raised one hand and flexed its fingers and wondered how this thing, this fleshy spider on the end of her arm, came to be hers, entirely at her command. She bent her finger and straightened it. The mystery was in the moment before it moved, the dividing instant between not moving and moving, when her intention took effect. It was like a wave breaking... There was no stitching, no seam, and yet she knew that behind the smooth continuous fabric was the real self—was it her soul?—which took the decision to begin movement and gave the final command.

—Ian McEwan, *Atonement*

“You,” your joys and your sorrows, your memories and your ambitions, your sense of personal identity and free will, are in fact no more than the behavior of a vast assembly of nerve cells and their associated molecules. Who you are is nothing but a pack of neurons.

—Francis Crick, *The Astonishing Hypothesis*

Nothing approaching the truth has yet been said on this subject.

—Thomas Nagel, *The View From Nowhere*

Up in the ivory heights, the free will debate has been raging for centuries—first in philosophy, and more recently, and perhaps more fiercely, in the brain sciences. It has been a divisive argument between the allegedly naïve position that people are in conscious control of their actions and the counterintuitive position that this experience of free will is illusory and people are automatons being pushed around by the compendium of known forces in a physical world.

For all the ostensible importance of such a question, the debate has had little to no impact outside of academia. For most people, the apparent volition behind their own behavior is satisfactory. What people want to happen, happens, and people generally assume that their conscious responses caused the outcome (Wegner, 2005). The rejoinder espoused by many scientists is that the connection between volition and action is merely correlational, and not causal. But this is held by many to be a purely academic argument that is at best beside the point and at worst absurd. Meanwhile, many philosophers of free will are satisfied with the compatibilist understanding that any type of free will worth wanting, such as freedom from coercion, compulsion, and political oppression, is unthreatened by scientific findings (Dennett, 2003). Although from time to time, people have wondered about whether this academic debate might have repercussions for the lay public (Breer, 1989; Skinner, 1971), such musings have been largely speculative, as no empirical evidence one way or the other has been brought to bear on the issue.

A recent set of studies, reviewed below, provides preliminary evidence that exposure to academic claims regarding the absence of free will can have an impact on moral action (Vohs & Schooler, 2006). Having one's traditional understanding of free will disturbed by the determinist argument seems to encourage a form of moral laxity. Contrary to the view that discussions of free will are largely academic, this work suggests that the belief in free will, be it justified or mistaken, affects behavior. Although it may be true that free will is an impotent epiphenomenon, the belief in free will can have real and potent consequences. Should the illusory free will position advanced in academic circles enjoy popular support among the lay public, it may be accompanied by larger social implications. The message that there is no free will may go from being understood as "nothing is controllable" to "everything is permitted." Again, regardless of the actual status of free will, a scientifically backed repudiation of it may encourage debauched behavior.

Put simply, the question of free will matters. And it matters not only to scientists in labs and philosophers in armchairs, but to the way that people live their lives. In a free society, neither scientists nor philosophers should be asked to suppress their views for fear of the possible social ramifications that the expression of such views might have. Nevertheless, if science is to be used as foundation upon which to promote claims that may have social impact, then the soundness of those claims deserves particular scrutiny. In this chapter, we consider the question of whether science has reached the point at which the

notion of free will must be dismissed, as some authors have argued (e.g., Crick, 1994; Pinker, 2002). Drawing on parallels between the challenges of conceptualizing free will and challenges in resolving the “hard problem” versus easy problems of consciousness (Chalmers, 1995), we argue that there exist two general classes of problems associated with the question of free will. The easy problems are those that have recently been posed within psychology—issues of automaticity, timing and backward referral, the neural systems involved, and so forth. These problems (reviewed in depth below) are by no means trivial, but they are akin to the easy problems of consciousness, which is to say that they are problems that can conceivably be solved using methods that are currently at the disposal of the scientific community. The hard problem of free will is different. Intricately related to the hard problem of consciousness, the hard problem of free will represents *the* core problem of conscious free will: Does conscious volition impact the material world? In other words, can phenomenal experiences translate into a physical events? And if so, how?

Although making this division should face similar criticisms as the ones leveled against Chalmers (e.g., Dennett, 1995) 10 years ago, there is particular need to do so. Without a clear distinction between the classes of problems, it is easy to confuse progress toward the easy problem as evidence pertaining to the hard problem. By clarifying which class of problems each scientist is investigating, confusion will be muted, and scientists can maintain as focus on the scope of their conclusions.

THE HARD PROBLEM OF CONSCIOUSNESS

Before distinguishing between the hard and easy problems of free will, it will be useful to first review this distinction as Chalmers (1995) applied it to consciousness. According to Chalmers, the easy problems of consciousness involve those that can be tackled with the standard methods of cognitive science and can be explained with computational and neural mechanisms. These include, but are not limited to, questions such as how attention modulates conscious experience, the difference between wakefulness and sleep, and the reportability of mental states. Although these problems are challenging, they are labeled as “easy” because they can be conceptualized within neurocognitive explanatory systems, and tested accordingly.

Chalmers separates these easy problems from the hard problem of consciousness—how the brain gives rise to the subjective experience that humans recognize as consciousness. This problem is hard because it cannot be conceptualized within standard neurocognitive explanatory systems. Current neurocognitive systems offer no principled way to distinguish between a system that has experience from one that does not. The elusiveness of subjectivity is illustrated by the hypothetical notion of zombies: individuals who are organized and behave just

like us but are devoid of subjective experience. Because subjective experience is (obviously) entirely subjective, there is no principled way to distinguish a zombie that claims to have subjective experience but does not, from a person who genuinely experiences consciousness. A similar problem occurs in determining whether nonhuman species experience consciousness. How can humans know what it is like to be a bat (Nagel, 1974), or, for that matter, a cockroach, or an amoeba, or a computer? Where do humans draw the line, and on what grounds? The hard problem of consciousness is hard because, in contrast to the easy problems, it is unclear what type of evidence would constitute a solution for it. Indeed, some have suggested that it is beyond the capacity of the human brain to even fathom a solution (McGinn, 1999).

Countless articles have been written about the hard problem of consciousness as articulated by Chalmers and his predecessors (e.g., Levine, 1983). Although authors disagree on the likeliness of solving the hard problem of consciousness (with perhaps a few exceptions), most everyone agrees it has yet to be solved and acknowledge that it poses uniquely challenging difficulties.

Schooler's Note on the Hard Problem of Consciousness¹

One way of illustrating how the hard problem of consciousness differs from all other problems in science is to consider the following thought experiment. Imagine that an elite group of scientists who have received every imaginable recognition and accolade told you that they had solved the hard problem of consciousness and had developed a technique that definitively discerned what possessed subjective experience and what did not. These scientists use their pioneering innovation on you and conclude that you do not in fact have consciousness... you just think you do. Would you accept their conclusion? You surely would accept any other conclusion such a group of scientists might offer. But in this case, I hazard, you would be absolutely certain they had done something wrong. Ultimately, unlike anything else one knows, when it comes to the existence of one's own subjective experience, one's own first person perspective a priori trumps even the most authoritatively imaginable scientific evidence.

Dennett (1995) has argued that there is no hard problem of consciousness because there is nothing particularly special about the first-person perspective. If I understand him (and I have to confess, despite my best efforts, I am far from sure that I do), his position is that there is ultimately nothing privileged about the first-person perspective that cannot be comparably extracted from the third-person perspective. Or put another way, any information that we believe we possess exclusively from the first-person perspective is ultimately suspect. Although I am quite sympathetic to the view that introspective evidence can be questionable (Schooler & Dougal, 1999; Schooler & Fiore, 1997; Schooler & Schreiber, 2004; Wilson & Schooler, 1991), it

seems to me there is at least one thing that we can only know from the first-person perspective, and that is that we are conscious. In a prior writing, we (Schooler & Schreiber, 2004) interpreted Dennett's dismissal of the significance of qualia as indicating that conscious experience is itself an illusion. Although we were in good company in this reading of his "qualia disqualified" chapter in Consciousness Explained (1991; see Searle, 1997). I have learned both from other writings of his (e.g., Dennett, 1997) and from personal communications that he does not deny the existence of subjective experience. The question that I still have not resolved, however, is how he knows that subjective experience exists, if it is not on the basis of his own first-person perspective.

Ultimately, it seems the unique knowledge afforded by subjective experience is precisely that which cannot be communicated to anyone else. Thus, the only way to argue for such knowledge is to appeal to individuals' own first-person experiences. It is tempting to suggest that those who deny the importance of qualia are somehow lacking the unique qualities of subjective experience that seem so self-evident to others of us. Having watched Dennett enjoy a fine meal with as much relish as anyone, I rather doubt this view, but the fact is I can't be sure...and therein lies the hard problem of consciousness.

THE EASY PROBLEMS OF FREE WILL

In our view, the distinction between the hard and easy problem of consciousness has important parallels to the problems associated with free will. Like consciousness, there are some questions about free will that seem to be straightforwardly addressable within the context of existing cognitive frameworks. As will be seen, issues such as the timing of the experience of will, the situations under which the experience of will arises, and situations in which these experiences are misguided can all be naturally studied and naturally understood within existing psychological/mechanistic frameworks. The hard problem of free will arises in reconciling the subjective experience of genuine personal choice with scientific claims that all actions are the necessary consequence of preexisting causes. Just as there seems to be no place to put subjective experience into the material formula, so, too, there seems no place for genuine choice. In the following discussion, we first consider some of the important work that has been done in addressing the easy problems of free will. Importantly, the work included in this category is far from trivial, and indeed in many cases has offered major insights into understanding the concomitants of an experience of will. Moreover, in contrast to the studies addressing the easy problem of consciousness, studies addressing the easy problems of free will do offer some constraints on the hard problem of free will, in the sense that they set some important boundary conditions on what people might potentially have control over. Nevertheless, as will be argued, such studies fail to put the proverbial nail in the coffin of free will,

as they neither establish nor disconfirm whether conscious thought can ever be responsible for action.

Serious scientific study of free will and conscious volition can be said to have begun with Libet's experiments on the timing of conscious choice. In order to locate the temporal position of the conscious willing of an action in the chain of events that led up to the performance of a voluntary action, Libet and his colleagues (1983) designed an apparatus to measure self reports of volition and compared them to the readiness potential—an EEG measured indicator of when the neural processes that give rise to the action begin. Libet and colleagues found that what they identified to be the volitional impulse to begin the action occurred around 350 ms *after* the readiness potential had begun. Without temporal primacy, it is very difficult to attribute causation to the volitional impulse. Further research (Libet, 1985, 1999) went on to demonstrate that a cognitive mechanism craftily reversed these temporal positions in the conscious mind. The actual time of the conscious impulse was “backward referred” about half a second so that it seemed to have occurred before the action began. What people were becoming conscious of was a false memory of having decided to act, rather than an actual desire.

Almost all of the works involved in the recent deluge of anti-free will arguments have referenced this study despite a steady stream of criticisms questioning the actual meanings of the time delay and backward referral (Gomes, 1999; Klein, 2002), the self-reported conscious desire (Dennett, 1995) and the “freedom” of the involved action (Levy & Byne, 2004). Nevertheless, Libet's experiments have stood the test of time and become the scientific spine of the anti-free will movement. Interestingly, however, Libet's own interpretation is considerably different from this prevailing view.

However Libet's work has been interpreted since, he refrained from concluding that his results could definitively inform the hard problem of free will. Instead of abandoning the possibility that free will exists, Libet has taken an approach less reliant on traditional materialist methods of explanation. He proposed that free will exists in the reduced capacity of a selection process—what he calls the veto clause. Although actions are generated unconsciously, he suggests there is a conscious control function that gives the option to allow or prevent that action from coming to fruition. Libet explains that because this veto power would represent the genuine type of conscious free will that can exist only unconnected to any physically determining forces (an “unmoved mover” in Clark's [1999] language), no directly observable material evidence should be expected. Instead, as this is a genuinely conscious and subjective power, science should look to introspection and the self-reports of his subjects to support the existence of the phenomenon (this approach is discussed further in the final section). It needs to be recognized that the veto clause alone cannot be taken as serious evidence informing the hard problem. Instead of addressing how the mental could affect the physical, Libet assumes it does and then creates

a mechanism that could exploit this assumption. Progress into the hard problem would require a discussion of the mechanisms by which the subjective power of the veto clause actually affects the physical neurons of the brain, how it actually moves the meat. Until this happens, the veto clause remains an interesting mechanism that evades the actual question that Libet wants to be addressing.

Shariff and Peterson (2005) have made another attempt to reconcile the time lag with at least an indirect form of conscious control. This “close-enough theory of free will” suggests that people’s actual actions are unconsciously initiated by well-learned schemata that link together object perception and associated motor actions. The conscious control over such actions is relegated to the switching of attention in object perception. By choosing to pay attention to various salient features of your environment, you elicit the relevant action schemas associated with that stimulus. Which schema will be initiated is probabilistically biased through a modulation of activation strength—so perceiving a glass of water as a “container of drinkable fluid” makes it highly likely that the motor schema associated with “grabbing and drinking from said container” is elicited. Though this is not direct free will, the authors suggest that it is something close enough. Their key insight is that instead of having complete control over one’s thoughts and behavior, people have some sort of detached and lower resolution control. An apt analogy may be to consider sailors who have no control over the wind and waves when sailing, nevertheless, they can set their sails and plot a general course.

Unlike Libet’s veto clause, Shariff and Peterson’s theory is embedded within and built out of a collection of converging clinical, cognitive, and theoretical support. Research on the peculiar neurological disorder of utilization behavior is synthesized with evidence from cognitive neuroscience as to how perceptual-motor “action macros” are construed and unconsciously run. A revised version of Gibson’s (1979) direct perception is interpreted in the context of modern theories on attention. Finally, the authors suggest how this close-enough theory can be reconciled with the timing issues introduced by Libet’s original data. But like Libet, the Shariff and Peterson theory does not tackle how conscious free will could exist. Instead, the theory proposes only how free will could work, with its existence assumed. Until this theory explains how conscious subjective selves can affect the redirection of attention in object perception, it does not address the central problem, or, in our terms, the hard problem.

Wegner’s (2002) argument faces the same limitations in addressing the hard problem of free will, but from the other camp. In *The Illusion of Conscious Will* (2002), he explores many examples of cases in which will is perceived, but is demonstrably uninvolved in action. Cases such as hypnosis, alien hand syndrome, and those involving direct brain stimulation suggest to Wegner that the neural systems governing the experience of will may be quite detached from the mechanisms responsible for the actions themselves. Moreover, Wegner details people’s tendency and perhaps motivation to perceive causal ownership to actions.

Wegner's explanations may be *necessary* to explain how a lack of conscious free will can be reconciled with the subjective experience of such volition, but it is not *sufficient* to dismiss such volition as an illusion. The existence of conscious free will is still fully compatible with the various illusions and mechanisms that Wegner describes. Like Libet, Wegner acknowledges the scope of his argument:

Questions of whether thought actually does cause action, for example, have been left in peace, and the issue of the role of consciousness in the causation of action has been ignored as well. This is because the focus of this theory is the experience of conscious will, not the operation of the will.... This theory is mute on whether thought does cause action. (Wegner, 2005, p. 32)

As has been demonstrated, neither Libet's data, nor the theories put forward by Libet and Shariff and Peterson, nor the mechanisms identified by Wegner address the core of the free will problem—the hard problem of free will. To use Wegner's words, they tiptoe “quietly around the big sleeping problem of real mental causation” (2005, p. 32). Each makes progress on surrounding issues, often in conflicting directions, but with regard to the hard problem, none does more than push it further back—shrinking the black box, but never opening it. Before we turn to some brave forays into that problem, it is worth looking at two more undertakings that could be considered “easy problems” of free will.

Bargh and his collaborators have assembled a wide-ranging body of literature that demonstrates the extent to which people's actions, emotions, and beliefs can be shaped unconsciously (see, e.g., Bargh & Chartrand, 1999). In one study, the researchers demonstrated that the mere mention of words relating to stereotypes of the elderly, such as “Florida” and “wrinkle,” caused a measurable and significant change in their walking speed. In another experiment, subjects primed with rudeness, as opposed to politeness, were more than four times as likely to interrupt a conversation between the experimenters (Bargh, Chen, & Burrows, 1996). A third showed that those subliminally primed with African American faces were more likely to behave aggressively in social interactions (an African American stereotype). Subsequent research has shown that unconscious priming can dramatically influence goal pursuit (Bargh, 1989), moral behavior (Shariff & Norenzayan, in press), self-sufficient behaviors (Vohs, Mead, & Goode, 2006) and, remarkably, the amount of free choice people think they had in a task (Seguin & Pelletier, 2000). In all of these studies, the subjects remained unaware of the priming, claiming full ownership over their actions, and sometimes vigorously disagreeing with the experimenter when informed of how they had been manipulated.

These are important studies. They compellingly illustrate not only the massive capabilities of the unconscious mind but, combined with Wegner's work, the disconnect between the feeling of free will and free will itself. People are *easily*

duped. At least some of what is perceived as spontaneous and endogenously generated choice is due to unnoticed environmental suggestion. There are, however, limits to what conclusions should be drawn from this research. Studying circumstances in which free will necessarily does not occur is not the same thing as demonstrating that free will can never occur. In many of the studies run by Bargh and colleagues, the standard procedure involves using an unconscious priming technique to observe the effects it has on some aspect of behavior. In each case, the participants' responses are, to a significant extent, the product of the experimenter's suggestion. Although keeping these suggestions outside of awareness is, of course, necessary for the purposes of the study, doing so restricts the opportunity to observe free will in the hard sense, or even the softer sense of a freedom of action. The priming technique is coercive and mere steps away from hypnosis. Using this research to directly inform the understanding of free will would be like using black-and-white photographs to study color blindness.

Moreover, like with visual illusions, illusions of will illustrate the manner in which subjective experience can be tricked, and thus offer important insights into sources of slippage in the system. However, to suggest that illusions of will demonstrate that the experience of will *never* has any bearing on what a person does would be akin to arguing that the existence of visual illusions demonstrates that there is never a correspondence between perception and external reality. Illusions of will importantly constrain the relationship between subjective experience and action, but they do not eliminate the possibility of such a relationship, and thus they leave the hard problem of free will intact.

THE HARD PROBLEM OF FREE WILL

The hard problem of free will boils down to whether the subjective experience of volition influences action, and if so, how? Like the hard problem of consciousness it involves understanding the mechanism underlying the interface between the mental and the physical. However, unlike consciousness, the existence of which is demonstrated by the very experience of subjectivity, the experience of free will does not in itself prove that it exists. Free will *could* be an illusion. This highlights a difference, then, between philosophical zombies and robots. Whereas a zombie is construed as a being that acts identically to a conscious human but lacks conscious experience, a robot, as we describe it, would still have those experiences but they wouldn't do anything. The robot would be driven entirely by its internal and unconscious machinery. Its consciousness would be an epiphenomenon—its perceived efficacy, an illusion. The most important difference is that people personally know, via introspection, that they are not zombies but they cannot know whether they are robots.

The similarities and differences between the hard problems of consciousness and free will require a reexamination of which methods of study can be trusted.

What role should introspective data play in understanding conscious free will? Libet, for instance, relies heavily on such introspective data to support the existence of his veto clause. Noting that consciousness, residing in a totally different ontological realm, is observable only subjectively, he explains that a mechanism involved in true conscious free will would exist in the same way—unapparent to direct objective observation but readily apparent phenomenologically. The veto mechanism fits this description, and as there is no clear evidence contrary to its existence; Libet argues on introspective grounds that it must exist.

Unsurprisingly, not everyone agrees. Clark (1999) lambastes such mentalist theories. “Feelings or intuitions,” he says, “never count as self-evident proof of anything” (p.286). But this is not wholly true. We reviewed one case, perhaps the only one (but certainly an important and relevant one), in which feelings and intuitions *do* count as proof. If the feeling of consciousness did not matter, if it contributed nothing toward proving that it exists, then science wouldn’t have any of these problems—easy, hard, whatever. This is where the role of introspective data, which is easily dismissed by the materialists, becomes thornier. Were one to take the line of reasoning that the materialists use to explain the infeasibility of conscious will and apply it to consciousness itself, one would have to conclude that there is no such thing as subjective experience at all. Were one to use, for consciousness, the line of reasoning that the materialists are using to explain the infeasibility of conscious will, one would quickly conclude that there was no such thing as subjective experience itself. But the *prima facie* introspective self-evidence of the existence of consciousness directly contradicts materialist arguments that it cannot exist. The fact that people’s feelings and intuitions conflict with the materialist argument at least casts suspicion on the comprehensiveness of materialism as it is currently understood. Consciousness manages to exist, in spite of the fact that from a material standpoint it should not. It is an anomaly that has yet to be satisfactorily accounted for. Hence, when considering conscious free will—disparaged by all the same argument that should make consciousness go away—one wonders if there isn’t room for another anomaly in this otherwise physical world.

Of course, speculating that free will is a plausible second anomaly in an otherwise materialist world must be done with great caution. Just because consciousness seems to trump materialist arguments on introspective grounds does not mean the floodgates should be opened for other nonmaterial constructs. Whereas each one of us would be willing to stake our very lives on the introspective certainty that we are conscious, perhaps none of us would be prepared to do the same for free will. Nevertheless, the two apparent phenomena seem to emerge from the same mystery—a misunderstood relationship between physical matter and subjective sensation. Given this connection and given the introspective power of the experience of personal agency, it should at least lead scientists to entertain the possibility that free will may similarly exist despite its material complications.

Ultimately, the landscape of the hard problem of free will may become clearer than it is now once there is a reasonably robust solution to the original hard problem of consciousness. Unfortunately, such a solution has been far from forthcoming—with some going so far as to suggest that such a solution is theoretically impossible to be grasped by the human brain (McGinn, 1999). Whether or not this position is true, it is worth pressing on and examining the way different theorists have at least attempted to address the problem. Although we refrain from articulating a single approach, we encourage the reader to keep an open mind as to the faults and strengths of the following positions, which we believe to be promising directions.

There are three broad camps in the battle for free will—hard determinist, compatibilist, and libertarian—and membership to these groups by and large falls along materialist-mentalist party lines.

Hard Determinism

The determinists are the most likely to dismiss free will as an illusion and reduce human beings to robots. They take the reductionist position suggested by Crick's quotation at the beginning of this chapter as their starting point. In the causal chain of behavior, they say, there is no room for anything beyond the story told by the physicists and neuroscientists. All behaviors and cognitions are initiated by the underlying "machinery" of the brain—whereas consciousness is another unidirectional product of these neural processes. It has no reciprocal effects on cognition, making conscious free will not only an illusion, but an impossibility. The brain simply does not work that way.

A quotation by Samuel Johnson that determinists are fond of using reads "all science is against free will, all experience is for it." This is absolutely true. With the exception of some research on the easy problems of free will (which *should*, as we've discussed, be excepted), there has been no evidence from traditional science that has provided a compelling case that free will in the sense of conscious agency exists. Experience certainly has, but as the determinists will be quick to point out, when pushed, experience alone has historically never been a good indicator of anything. Whether consciousness, and subsequently conscious free will, break this historical mold and should be given exceptional status in the realm of science depends on just how special and unique a phenomenon one thinks consciousness is, which is a matter of no small disagreement between researchers in the field.

Finally, it is important to note that when it comes to consciousness and free will, impotent does not mean the same thing as useless. This is a crucial distinction that is often lost in the imprecision of our lexicon. Although conscious free will may not exist, this does not mean that the *idea* of conscious free will does not exist. The belief in free will exists in the brain, and, as the Vohs and Schooler

data show, the belief in free will is involved in the causal chain of behavior. If it is an illusion, the illusion is an efficacious one and insofar as it prevents moral laxity, it is a *useful* one. It keeps people feeling morally responsible as agents, although, in truth, they may not be. There may be a time, if the hard determinists are proven right, that people will choose to maintain a belief in free will that they know to be disingenuous—acting as if free will still existed and deceiving themselves into good behavior. But we are getting ahead of ourselves; the determinist position has not been proven, nor has it been disproven. It has, however, been challenged extensively, and these other positions merit consideration.

Compatibilism

The term compatibilism is used generally and originally with reference to the debate within philosophy about free will in the context of an allegedly physically determined universe—a debate that we have stayed away from. Therein, it referred to the position that free will could coexist with—be *compatible* with—a deterministic view of the universe. We modify this meaning only slightly to refer to psychological compatibilism that is, as one would expect, the position that attempts to reconcile free will with the apparent impotence of consciousness. Though it is a broad category, most psychological compatibilists tend to eschew appeals to material dualism without completely dumbing humans down to robots. They maintain, in one way or another, that despite the deterministic laws people must follow, they nonetheless have genuine options and opportunities for action.

Certain compatibilist interpretations are, no doubt, less compelling than others. For instance, some seem to be saying that so long as people do not feel as though their choices are psychologically or physically constrained, then they enjoy free will (Strawson, 1998). In the light of both Bargh's (e.g., Bargh & Chartrand, 1999) and Wegner's (2002) research, this position falls prey to fatal weaknesses. The feeling of free will, as Bargh and Wegner have shown, is an easily fooled indicator of true free will—even in the sense that these compatibilists mean. Bargh's subjects in numerous experiments were undoubtedly psychologically compelled to behave in the way they did, but felt throughout that they were exercising their free choice in all their endeavors. Feeling, we must repeat, is not enough. It is an idea worth considering that the *belief* in free will may be the only "variety of free will" worth wanting, but this is a separate debate. Ontologically, this form of compatibilism is left wanting.

A more sophisticated compatibilist approach stems from Velmans' *dual-aspect* approach. Dual-aspect theory is an attempt to frame the hard problem of consciousness by employing a softer form of Cartesian dualism. Velmans and others square a materialist understanding of the world with the apparent differences between third-person neurons and first-person experience by arguing

that neurons and experience can be made of the same materials, but still have different identities. Velmans (2002), specifically, explains that although the two identities share the same “informational structure,” they are ontologically different. This type of theory, which shows first- and third-person experience to be *equivalent*, but not *identical*, is known as identity dualism.

Relating identity dualism to free will provides a potentially promising theory of volition. Clark (1999) outlines such a position, an informed synthesis of both Gomes (1999) and Claxton (1999) from the same volume. Moving from the understanding that the subjective mind *is* brain-based activity, he explains that conscious desire doesn’t spawn or lead to neural processes any more than neuronal activity spawns or leads to conscious experience. There is no two-step process. Instead, the two are equivalent things; the experience of conscious free will is the first-person perspective of the neural correlates of choosing. This means that although there is no nonphysical, conscious self-homunculus that sets things in motion, there is still an “I”—the holistic compendium of your brain processes over which people can and should take ownership. It is this holistic self with which people should identify. At the risk of confusing the issue, one could say there is no prime mover, but rather the movement is the mover.

In some ways, being compelled by compatibilist approaches like this one leaves a person with the feeling that he or she had a run-in with a known pick-pocket: One should probably do a quick pat-down to see if anything important was taken. When considering the Velmans-Clark-Gomes-Claxton approach, it is worth asking what aspects of free will are maintained, what are lost, and what, if any, are gained. Ultimately, people maintain most of the types of free will (worth wanting) that they had when assuming that they possessed a nonphysical, conscious controller. They remain responsible for their actions; and they remain in charge. What they want still happens.

Nonetheless, there are subtle but deceptively important differences between the folk understanding of free will and this current reconceptualization. In the former position, one’s “I” is understood to refer to “the me that does thing thinking” and this self is credited with being the one that consciously controls one’s actions. The new approach dissolves the conscious self into the larger “I” and “the me that does the thinking” is embedded within the whole brain. “I” still control my actions, but the “I” is reconceived to be the coalition of my brain processes. And although I am still fully in control of my actions—although my will is still free—the nature of conscious free will changes. Instead of saying that people have conscious free will, we must instead say that people are conscious *of* their free will. Instead of saying that my consciousness (me) is making the decisions, we need to say that I am conscious of the parts of my brain (still me) that are making the decisions. Instead of saying the “I” moves the machinery of my brain, “I” *am* the machinery of my brain, and “I” consequently move myself.

Although this distinction may initially seem to be just semantic—rhetorical sleight of hand—closer inspection yields important differences, both

philosophical and pragmatic. Most importantly, it seems that one loses the ultimate sense of free will that Kane (1996), and many others, believe worth wanting. Giving up control from the conscious, agentic self to the holistic “you are your brain” controller means just that. The “I” that one usually identifies with—one’s conscious self—loses its claim on causation. Insofar as people continue to identify with the former conscious agent, they do not have free will in the way that Kane defines it: “*the power of agents to be the ultimate creators (or originators) and sustainers of their own ends and purposes*” (p. 4). To say, as the proponents of the new position will, that people maintain this free will by expanding their conception of themselves as agents to their holistic brain is fine and by a certain interpretation true, but it introduces a host of its own problems.²

Libertarianism

Libertarians, like the hard determinists, are incompatibilists. However, unlike the determinists, they believe that human beings still maintain a sense of ultimate agency and control over their own will.

The libertarian position is both the one that is most aligned with traditional, intuitive understanding of free will and the one that is most often, and most vociferously, maligned. So out of vogue is this position with the majority of consciousness researchers that Wegner, with only some irony, divided the field into “robogeeks” (those who espouse something approaching the determinist position) and “bad scientists” (those who stick to their intuitive beliefs despite the wealth of conflicting evidence). And the libertarians have been criticized, with some justification, as resorting to increasingly “panicky metaphysics” to maintain their belief in ultimate agency. By clinging to traditional conceptualizations of free will, they have been accused of resisting the progress of science and being unwilling or unable to update to new paradigms. This may be true to some extent. But this argument could also be made in the opposite direction: The assumptions involved in scientific examination may themselves be in error. The existence of consciousness and the apparent existence of conscious volition may be examples of anomalies that indicate the limits of our current investigative paradigm. And those researchers who are strictly abiding by the established materialist modes of investigation may be the ones who are being overly rigid, trying vainly to cram ever more complex phenomena into inadequate methods of explanation. It is perhaps not the traditional understanding of free will that is in error but, rather, the traditional understanding of how to do science.

Libertarians point, not unreasonably, to increasingly bizarre discoveries in other scientific endeavors such as quantum indeterminacy, discoveries that have forced scientists to reevaluate their preexisting assumptions. Coupled with the relative ignorance of the brain sciences, libertarians insist that there is ample room left for the claims of determinists of any persuasion to be overturned.

These claims are all certainly true, but it is this attitude that proves to be the greatest weakness of the libertarian position. From the start, the libertarians have been playing defense—trying to maintain a status quo position against a tide of criticisms. The assumption has been that the burden of proof was on those attempting to dispel the existence of free will. This has not been enough. For the libertarian position to be taken seriously, its adherents need to demonstrate its viability, and not just its possibility.

Schooler's Reflections on the Hard Problem of Free Will

In principle, the compatibilist perspective is ideal as it allows us to have our material cake while freely choosing to eat it, too. The problem is that I, for one, simply cannot get my head around compatibilism. I understand that determinism does not rule out the opportunity to make deliberate decisions that are free from coercion, I acknowledge that I can redescribe the control of my actions as being completed by my brain as opposed to my mind, and I admire Dennett's (2003) argument that the evolution of culture and the human brain's capacity for rational analysis has enabled individuals to make reasoned decisions. I just don't understand the following: If any noncoerced reasoned decision that I am about to make must necessarily be carried out in a specific manner based on a preexisting causal chain, then how can I be free to choose otherwise? And if I really have no option but to do exactly what I end up doing, how can it be said that my choice was free? On this point, I find myself in agreement with Greene and Cohen (2004) who argue, "...contrary to legal and philosophical orthodoxy, determinism really does threaten free will and responsibility as we intuitively understand them." (p.1780)

Given the indefensibility of compatibilism, Greene and Cohen (2004) go on to argue that neuroscientific evidence supports the hard determinist view that "free will as we ordinarily understand it is an illusion generated by our cognitive architecture." (p.1784) They, like many others, dismiss libertarianism as "panicky metaphysics" noting, "...there is not a shred of scientific evidence to support the existence of causally effective processes in the mind or brain that violate the laws of physics." (p. 1777). Importantly, however, an absence of evidence is not necessarily evidence of an absence. For example, there is no scientific evidence indicating the existence of other universes; nevertheless, many rational individuals have postulated that alternative universes might exist (Greene, 2003). The problem is that just as there is no evidence that alternative universes do exist, there is also no evidence that they don't. Similarly, although we have made important strides in understanding human behavior, we are still very far from perfectly predicting human action. Given the current indeterminacy of human behavior, there is still room for additional causal sources. Thus although there may be no scientific evidence for the existence of free will, there is also no scientific evidence of its absence.

Moreover, and here I go out even further on the limb, there actually is some evidence that consciousness can influence the outcome of physical events in a manner that, if true, would appear to violate at least our current understanding of physics. Specifically, a recent review (Bösch, Steinkamp, & Boller, 2006) in the prestigious journal *Psychological Bulletin* reported the results of a meta-analysis reviewing 380 studies that examined whether “random number generator output correlated with human intentions” (e.g., whether participants’ efforts to will a random number generator to produce odd numbers increased the likelihood of it doing so). The result of this analysis revealed a “significant but very small overall effect size.” (p.497) The authors cautiously concluded that “publication bias appears to be the easiest and most encompassing explanation for the primary findings of the meta-analysis.” However, they also conceded that this explanation would require a rather large number ($n = 1500$) of unpublished studies. Moreover, they acknowledged that the alternative possibility that human intention was influencing random events could not be ruled out, noting, “The effect in general, even if incredibly small, is of great fundamental importance—if genuine.” (p. 517)

Clearly, remarkable claims require remarkable evidence, and this report, though striking, is unlikely to persuade many that they should accept the claim that human consciousness can influence physical events. At the same time, it does constitute at least a “shred of evidence” that causally effective processes in the mind can influence physical processes. And if conscious will is capable of influencing random number generators at a distance, then surely it should be capable of influencing the far more proximal behaviors of the brain.

It is important to emphasize that even if mental intentions can influence physical events, it does not necessarily follow that free will exists. The human intentions that potentially influence random number generator could themselves be the necessary consequence of causal chains. My argument is simply that given the existence of at least a shred of evidence for causal effects of human consciousness, and given the degree to which the sources of human thought and action are still not understood, it seems premature to conclude that scientific evidence definitively rules out the possibility of a genuine impact of conscious will.

Ultimately, the viability of the libertarian perspective will depend on the generation of accounts that are both phenomenologically compelling and scientifically tractable. In this regard, I am sympathetic to the view that quantum physics may offer a possible opening for genuine free will. There are a number of characteristics of quantum physics that offer some hope for salvaging of free will. First, quantum physics demonstrates that causal determinism is not, as is often suggested, a necessary aspect of the universe that is fundamentally violated by the notion of free will. If, as quantum indeterminacy suggests, the future is not written in stone, then different actions could lead to different futures. It is often pointed out that quantum indeterminacy offers little solace for libertarians, because having one’s choices influenced by a combination of deterministic forces and some random quantum element still leaves no room for the conscious chooser. As Greene and Cohen (2004) observe, “If it

turns out that your ordering soup is completely determined by the laws of physics, the state of the universe 10,000 years ago, and the outcomes of myriad subatomic coin flips, your appetizer is no more freely chosen than before" (p. 1777). However, this presumes that the indeterminacy observed at the quantum level is entirely random. If, as the studies on random number generators mentioned above hint, the mind can influence the outcome of random events, then indeterminacy associated with quantum effects might appear random only because we have failed to assess the causal impact of consciousness.

A second important feature of quantum indeterminacy is its peculiar relationship with observation. Although accounts of the role of observation in quantum outcomes remains an area of considerable dispute, many reasonable scientists have argued that the way in which we observe quantum events influences how they unfold (for a fascinating and highly readable account of the challenges of understanding the relationship between physics and consciousness written by noted physicists, see Rosenblum & Kutner, 2005). If consciousness can influence external events, then why couldn't it similarly influence internal mental events? Admittedly, at present, it is not clear how quantum indeterminacy could operate at the macro level of warm brains. Although some have speculated about possible mechanisms (e.g., Hameroff & Penrose, 1995), others have highlighted the various reasons that such accounts are implausible (Koch & Hepp, 2006).

Such disputes lead to the final important lesson of quantum physics, which is that although its properties can be described with remarkable precision, the explanation for why it interacts with observation in the way that it does remains a central mystery. Given the further inability of science to account for consciousness more generally, it seems that we simply do not understand the nature of either physics or consciousness enough to know with assuredness what their relationship might be. Were subjectivity and volition modest aspects of our existence, then it would seem entirely unreasonable to suggest that some unforeseen revolution in scientific understanding might offer them a greater place. But to the contrary, the experience of being and doing are arguably the two most essential aspects of our day-to-day lives. Although science has made incredible advances, has it really progressed to the point at which it can reasonably ask us to disregard the defining aspects of our existence?

THE CONSEQUENCES OF A BELIEF IN FREE WILL

Although the above approaches to the hard problem of free will ultimately be judged on their scientific and philosophical merits, there are undoubtedly other factors affecting the attractiveness of each position. The libertarian position, for instance, has been vigorously defended in large part because of phenomenological factors. There is the intuitive appeal of the mind working the way one feels it does. There is also a sentimental attachment to the belief that as human beings, we are each prime movers. Perhaps as a consequence, there is a visceral

repulsion to the determinist doctrine of fully automated robots. Agency and responsibility are tied to meaning, and being robbed of agency likewise robs one of meaning (Heine, Proulx, & Vohs, 2006). The compatibilists, meanwhile, seem to be resigning themselves to the indefensibility of the libertarian position, but trying to salvage as much meaning as they can through complex arguments.

Beyond the psychological palatability afforded to each position, there is also the moral component. The question of free will has always been tied to the moral implications of the argument (Pereboom, 1997). Generally, the predictions made about the consequences of an eroding sense of free will differ systematically with the position they are trying to advance.

The libertarians tend to present morally bankrupt dystopias as the inevitable consequence of abandoning a belief in ultimate agency and responsibility. The argument, an old one, can be summed up as follows: Without belief in causally responsible agents who could have done otherwise, there is nothing to blame or praise and therefore no way to sanction moral behavior. Without such sanctions, without the social and personal prescriptions that come with recognizing what is worth praise or blame, and without the personal dignity that comes with agency, people will be reduced to selfish beings without a moral compass. This position has been recently articulated by Bennett (1998) and Goodwin (1998), the latter of whom believes it to be held by the lay majority.

Determinists and compatibilists present rosier views. Greene and Cohen (2004) predict that as people become more accepting of compatibilism in the justice system, they will move from a punitive system demanding retribution to a more humane system more concerned with the consequences of punishment. Clark (1999) agrees, adding that by replacing the latently metaphysical beliefs that most laypeople hold with "a thoroughly naturalistic conception of the self and its choices" (p. 17), society may experience a less punitive culture more focused on the exogenous causes of individual ills, whereas individuals, themselves, will benefit from a lack of self-consciousness and personal blame. This optimistic outlook also includes less wealth hoarding as ambition softens and, with a shift in societal understanding of freedom, a "more responsible use of such freedom" (p. 18).

Unfortunately, empirical investigation from our lab suggests otherwise. People who are disabused of the illusion of agentic control seem to, at least temporarily, abandon their moral code. Two experiments that manipulated participants' belief in free will show that when people come to believe that the idea of free will is untenable, they behave amorally, for which we found evidence in the realm of cheating (Vohs & Schooler, in press).

In the first experiment, participants were induced to believe or not believe in free will via reading an essay written by Nobel laureate (given for the codiscovery of DNA) Francis Crick. One chapter of Crick's book, *The Astonishing Hypothesis* (1994), claimed that rational, thinking people (such as scientists) long have denounced the idea of free will, noting that it is instead

a byproduct of the human mind. Participants who read this essay were in the anti-free will condition, whereas control condition participants read from another chapter in this same book on consciousness, which did not contain references to free will. Afterward, participants were seated in front of a computer that displayed, one by one, multicomponent mathematical problems that participants were supposed to calculate in their heads. We told them that due to a programming error, the computer had a glitch that allowed the answer to be shown after a short period of time. Participants were told that they could, however, stop the answer from being shown by pressing a certain key on the keyboard. This situation, then, gave participants the opportunity to cheat, but also gave them a simple way to avoid it—a slight movement of the hand. Would convincing participants not to believe in free will alter whether they let themselves cheat?

As expected, it did. Participants who were disabused of the idea of free will cheated more by letting the answer appear relative to participants who read an essay that did not speak to the existence of free will (Vohs & Schooler, in press). The total number of math problems that participants saw was 20, and the anti-free will group let themselves cheat on almost 12 problems, whereas the control group let themselves cheat 9 times.

Some readers may be wondering whether the Crick essay effectively changed participants' beliefs about free will, or whether the effects were due to some other cause. To test for changes in cognitions, we included a scale that measured belief in free will (Paulhus & Margesson, 1994). As expected, participants in the anti-free will condition reported lower scores on the free will subscale compared to participants in the control condition. Moreover, belief in free will scores correlated significantly (and negatively) with the propensity to cheat. In other words, after reading the anti-free will essay, compared to measures taken after reading an essay that was devoid of free will information, participants said they were unconvinced that free will exists and they let themselves cheat.

Although these results were promising, we recognized that there are alternate explanations of these results. It may be that although cheating behavior changed, participants were not acting immorally but rather behaving passively as a result of the anti-free will essay. Hence, we conducted another experiment to rule out this explanation, using a cheating behavior that required active participation to cheat. We also added a condition that enhanced beliefs in free will, which provided an opportunity to test the full effects of believing or denying the existence of free will.

In this second experiment (Vohs & Schooler, in press), participants received one of three treatments. In one condition, we had participants read a series of statements designed to induce a feeling of determinism, which we believed would also have the effect of reducing free will. Sample statements included, "Ultimately, we are biological computers—designed by evolution,

built through genetics, and programmed by the environment.” Participants’ task was to read each statement and think about it, and then when instructed they were to turn the page and read another statement. This task is modeled after the oft-used Velten mood induction task (Velten, 1968). In another condition, participants read statements that were designed to bolster beliefs in free will, such as “I am able to override the genetic and environmental factors that sometimes influence my behavior.” A third group of participants read neutral statements.

The cheating opportunity was set up such that participants self-scored a cognitive test on which they were to be paid \$1. Ostensibly because of an unexpected errand, the experimenter left the room and allowed participants to score their exam and then pay themselves for their performance on the test. We compared the money participants paid themselves, as a proxy for their claimed scores on the exam, to veridical scores from participants who took the exam and were not allowed to self-score. The research question was whether participants would give themselves differential amounts of money as a function of whether they had been induced to belief in free will, determinism, or whether their beliefs were left unchanged.

The results showed that after participants read statements that told them their actions were predetermined and therefore not under their control, they cheated more—as evidenced by more money taken in this condition compared to the control condition and the free will condition. Reading statements that bolstered participants’ belief in free will did not affect cheating behavior, as these participants paid themselves as much money as did participants whose scores were known. Once again, we knew that participants’ beliefs did change, as evidenced by changes in the Free Will and Determinism Scale (Paulhus & Margesson, 1994) as a function of condition.

Hence, telling people that free will does not exist or telling them that their behavior is caused by predetermined mechanisms outside of their control leads them to cheat more so than people who are not induced to change their beliefs about free will. Note, interestingly, that in the second experiment, the control condition and the free will bolstering condition were not significantly different from each other, which suggests that lay beliefs about free will are in line with ideas of ownership and authorship of one’s own behavior (cf. Wegner, 2005).

These data lend empirical support to the suspicion that eliminating people’s beliefs in free will may be accompanied by reduced ethical behavior. These data do not mean, of course, that scientists should put all the weight behind the libertarian position. Nor does it mean that society will, in the long term, never be able to morally adjust to compatibilist³ or even determinist worldviews. Instead, these data suggest that there are real consequences to scientists’ claims. There may come a time when the evidence will require society to rethink its conceptions of freedom and responsibility—as Dennett says, “institutions and practices based on obvious falsehoods are too brittle to trust” (2003, p. 290). But

the key word in that statement is *obvious*. Nothing, yet, is obvious. Given the present lack of a solution to the hard problem of free will and given the societal repercussions of convincing people that they lack genuine control of their behavior, it seems that caution is warranted when making assertions regarding free will outside of the ivory tower.

NOTES

1. In several places in this chapter, we have agreed to break with the convention of writing co-authored papers with an exclusively unified voice. The metaphysical ramifications of free will and consciousness are particularly contentious, and one of us (Schooler) has controversial opinions on the matter that are not necessarily shared by the other authors. As a compromise, Schooler has inserted several brief reflections sections (set apart from the main text) into this chapter that should not be assumed to reflect the views of Shariff and Vohs.

2. Some of these problems are particularly evident in situations in which the system has broken down. Consider, for example, alien hand syndrome, wherein one's limb acquires a "mind of its own" often interfering with willed activities and even attempting to choke the "host" (Scepkowski & Cronin-Golomb, 2003). In this case, the sense of control that we are conscious over is overwhelmed by the control of another facet of the brain. The brain has divided itself into two competing agents. It would be hard to convince the patient to identify with his or her holistic brain when unconscious parts are rebelling against the conscious parts. In this case, it is hard to see the pragmatic difference between having alien hand syndrome, in which the hand is controlled by endogenous factors, and having a condition in which one's hand has actually come under the control of exogenous aliens. To call this free will is surely to strip the term of any meaning. Although this example and others like it (e.g., schizophrenic delusions of alien control or more commonplace diseases such as Parkinson's disease) represent clinical situations in which the mechanisms of will and ownership have gone awry, they are illuminating and worth considering.

It isn't only in clinical cases that such issues arise, however. We have to consider, for example, what identification with the holistic self means for the unconscious or nonconscious decisions that are made by our brain. What about cases of sleepwalking, or reflexes or decisions induced by Bargh's experiments? These issues demonstrate, again, that adopting compatibilist positions will require people to sacrifice at least some of the intuitive, libertarian position that most have grown comfortable with. Moreover, these sacrifices might not be immediately apparent or easy to understand.

3. One important future direction is to examine the impact on moral behavior of exposure to a compatibilist perspective. Dennett (this volume) has suggested a description of a compatibilist worldview that he argues might not only avoid inducing immoral behavior, but might actually facilitate the moral behavior of participants. Notably, however, his suggestion of how to characterize the compatibilist worldview does not explain compatibilism; it simply assures the reader

that philosophers have worked it out. Thus the potential efficacy of his suggested manipulation may simply take advantage of people's willingness to accept claims on the basis of authority.

REFERENCES

- Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 3–51). New York: Guilford Press.
- Bargh, J. A., & Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist*, 54, 462–79.
- Bargh, J. A., Chen, M., & Burrows, L. (1996). The automaticity of social behavior: Direct effects of trait concept and stereotype activation on action. *Journal of Personality and Social Psychology*, 71, 230–44.
- Bennett, W. (1998). *Defining the human spirit: Historical perspectives and contemporary challenges*. Presented at Neuroscience and the Human Spirit, Washington, D.C. September 24–25.
- Bösch, H., Steinkamp, F., & Boller, E. (2006). Examining psychokinesis: The interaction of human intention with random number generators—a meta-analysis. *Psychological Bulletin*, 132, 497–523.
- Breer, P. (1989). *The spontaneous self: Viable alternatives to free will*. Cambridge, MA: Institute for Naturalistic Philosophy.
- Chalmers, D. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219.
- Clark, T. W. (1999). Fear of mechanism: A compatibilist critique of the “The Volitional Brain.” *Journal of Consciousness Studies*, 6(8–9), 279–93
- Claxton, G. (1999). Whodunnit? Unpicking the “seems” of free will. *Journal of Consciousness Studies*, 6, 99–114.
- Crick, F. (1994). *The astonishing hypothesis*. New York: Scribner's.
- Dennett, D. C. (1991). *Consciousness explained*. Boston: Little, Brown & Co.
- Dennett, D. C. (1995). Facing backwards on the problem of consciousness. *Journal of Consciousness Studies*, 3(1), 4–6.
- Dennett, D. C. (1997). *Kinds of minds: Towards an understanding of consciousness*. New York: Basic Books.
- Dennett, D. C. (2003). *Freedom evolves*. London: Allen Lane.
- Gibson, J. J. (1979). *Ecological approach to visual perception*. Boston: Houghton Mifflin.
- Gomes, G. (1999). Volition and the readiness potential. *Journal of Consciousness Studies*, 6, 59–76.
- Goodwin, F. K. (1998, September 24–25). *Opening remarks*. Presented at Neuroscience and the Human Spirit, Washington, DC.
- Greene, B. (2003). *The Elegant Universe: Superstrings, Hidden Dimensions, and the Quest for the Ultimate Theory*. New York: W.W. Norton and Company.
- Greene, J. D., & Cohen J. D. (2004). For the law, neuroscience changes nothing and everything. *Philosophical Transactions of the Royal Society of London B*, 359, 1775–85.

- Hameroff, S. R., & Penrose, R. (1995). Orchestrated reduction of quantum coherence in brain microtubules: A model for consciousness. *Neural Network World*, 5(5), 793–804.
- Heine, S. J., Proulx, T., & Vohs, K. D. (2006). Meaning maintenance model: On the coherence of social motivations. *Personality and Social Psychology Review*, 10, 88–110.
- Kane, R. (1996). *The significance of free will*. New York: Oxford University Press.
- Klein, S. (2002). Libet's temporal anomalies: A reassessment of the data. *Consciousness and Cognition*, 11, 198–214.
- Koch, C., & Hepp, K. (2006). Quantum mechanics in the brain. *Nature*, 440, 611–12.
- Levine, J. (1983). Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly*, 64(4), 354–61.
- Levy, N., & Byne, T. (2004). A will of one's own: Consciousness, control, and character. *International Journal of Law and Psychiatry*, 27, 459–70.
- Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, 8, 529–66.
- Libet, B. (1999). Do we have free will? *Journal of Consciousness Studies*, 6(8–9), 47–57.
- Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness potential): The unconscious initiation of a freely voluntary act. *Brain*, 106, 623–42.
- McEwan, I. (2001). *Atonement: A novel*. New York: Doubleday.
- McGinn, C. (1999). *The mysterious flame: Conscious minds in a material world*. New York: Basic Books.
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83(4):435–50.
- Nagel, T. (1986). *The view from nowhere*. New York: Oxford University Press.
- Paulhus, D. L., & Margesson, A. (1994). Free Will and Determinism (FAD) scale. Unpublished questionnaire, University of British Columbia, Vancouver, Canada.
- Pereboom, D. (1997). *Free will*. Indianapolis, IN: Hackett.
- Pinker, S. (2002). *The blank slate*. New York: Viking.
- Rosenblum, B., & Kultner, F. (2006) Quantum Enigma: Physics Encounters Consciousness. New York: Oxford University Press.
- Scepkowski, L. A., & Cronin-Golomb, A. (2003). The alien hand: Cases, categorizations, and anatomical correlates. *Behavioral Cognitive Neuroscience Review*, 2, 261–77.
- Schooler, J. W. & Dougal, S. (1999). The symbiosis of subjective and experimental approaches to intuition. *Journal of Consciousness Studies*, 6, 280–87.
- Schooler, J. W., & Fiore, S. M. (1997). Consciousness and the limits of language: You can't always say what you think or think what you say. In J. D. Cohen & J. W. Schooler (Eds.), *Scientific approaches to consciousness* (pp. 241–57). Mahwah, NJ: Erlbaum.
- Schooler, J. W., & Schreiber, C. A. (2004). Experience, meta-consciousness, and the paradox of introspection. *Journal of Consciousness Studies*, 11, 17–39.
- Searle, J. (1997). *The mystery of consciousness*. New York: Review Press.
- Seguín, C., & Pelletier, L. G. (2000). *Automatic activation of intrinsic and extrinsic motivation*. Manuscript submitted for publication, University of Ottawa.

- Shariff, A. F., & Norenzayan, A. (in press). God is watching you: Supernatural agent concepts increase prosocial behavior in an anonymous economic game. *Psychological Science*.
- Shariff, A. F., & Peterson, J. B. (2005). Anticipatory consciousness, Libet's veto and a close enough theory of free will. In R. Ellis & N. Newton (Eds.), *Consciousness and emotion: Agency, conscious choice and selective perception* (pp. 197–215). Amsterdam: John Benjamins.
- Skinner, B. F. (1971). *Beyond freedom and dignity*. New York : Knopf.
- Strawson, G. (1998). Free will. In E. Craig (Ed.), *Routledge encyclopedia of philosophy*. Retrieved August 11, 2005, from <http://www.rep.routledge.com/article/V014SECT1>
- Velmans, M. (2002). How could conscious experiences affect brains? *The Journal of Consciousness Studies*, 9(11), 3–29.
- Velten, E. (1968). A laboratory task for induction of mood states. *Behaviour Research and Therapy*, 6, 473–82.
- Vohs, K. D., Mead, N. L., & Goode, M. R. (2006). The psychological consequences of money. *Science*, 314, 1154–56
- Vohs, K.D. & Schooler, J.W. (in press) The value of believing in free will: Encouraging a belief in determinism increases cheating. *Psychological Science*
- Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.
- Wegner, D. M. (2005). Who is the controller of controlled processes? In R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 19–36). New York: Oxford University Press.
- Wilson, T. D., & Schooler, J. W. (1991). Thinking too much: Introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology*, 60(2), 181–92.