8 Agency and Control: The Subcortical Role in Good Decisions

Patricia S. Churchland and Christopher L. Suhler

If nonconscious brain processes contribute to decision making, what difference, if any, should that make to our traditional conception of what it is to be in control and to be a responsible agent? This is the question that motivates our exploration below. First, a preliminary clarification.

"Free will" is an expression festooned with semantic bear traps. We prefer to avoid those. One assumption that frequently ensnares the unwary is that actions can be sorted into one or the other of two separate bins: freely chosen or not. In fact, however, actual decision making is far messier. Sometimes a person is very sleepy or very hungry or chronically stressed or desperately frightened (Arnsten, 2009). Sometimes a person may suffer brain damage. The neat two-bin model is utterly inadequate to the reality of decision making. Another hobbling assumption is that free will implies a freedom from *all* causal antecedents to the decision. No one can seriously maintain free will in this sense. Denying free will for such "must-be-uncaused" reasons sounds like you are denying what we all know to be a fact: Many people act in a controlled manner much of their lives.[1] Semantic bear traps galore.

The capacity for self-control, as is evident from research in psychology and neuroscience, is linked to certain causal antecedents and the functioning of specific neuronal structures and pathways. Moreover, self-control comes in degrees. It develops as the infant matures and can decline as dementia destroys. Self-control can be affected by factors such as stress, hunger, cold, and exhaustion. Without belaboring the point, we find that these considerations motivate a shift from the language of *free will* to the language of *control*.[2] Wrangling over the metaphysical esoterica of *free will* is apt to be unproductive, and in any case the metaphysical matters have been well discussed elsewhere (Churchland, 2002; Flanagan, 2003; Dennett, 2004).

Control we provisionally define as the capacity of an individual to act in an intelligent and adaptive manner within a particular environment—to maintain a goal, to defer gratification, and to suppress disadvantageous impulses. Control is crucial to matters in the legal domain, as well as to debates in philosophy concerning responsibility. It is also important to a more general set of assumptions regarding how to think about oneself.

Because *control* cannot at this point be precisely defined, it is also useful to amplify the provisional definition by noting the prototypical cases where there is agreement about application of the term (Johnson, 1993). In the prototypical case of controlled action (agency), a healthy adult human who is awake and cognitively unimpaired is said to be an agent when he walks into a bakery and buys a loaf of bread. His action is under his control. By contrast, a man who is sleepwalking and kicks his wife is considered not fully in control. A person who is a chronic nicotine addict has less than full control over his choice to buy a package of cigarettes though he is still held responsible for that action. A person who emits a startle response to a loud and unexpected noise, kicking over a lantern and thereby setting the barn on fire, has compromised control, and he is not held responsible. The prototypes are where agreement is maximal. The difficult cases arise when the person is not yet mature, is not fully awake, is ignorant on some crucial point, or is suffering from brain damage, addiction, or a psychiatric condition. Sometimes there may be no right answer as to whether a person was in control of his or her actions, though in a legal context a definite answer may be required nonetheless. Control, as with many other concepts, has a radial structure with declining degrees of similarity from the central cases to those in the fuzzy boundary.[3]

Recently, however, a growing body of research on the prevalence of nonconscious cognition and the influence of nonconscious factors on behavior has been taken to imply that even the prototypical cases of agency are not what they seem (Wilson, 2002). How can a person have any control at all over nonconscious factors? How can such factors enter into a person's conscious deliberations and reflections? How can they be part of a person's weighing of reason and evidence? And if you cannot have control over nonconscious factors, cannot make them part of your conscious deliberations, how can you be responsible for any action whatsoever? As some critics have put it to us in conversation, nonconscious cognition and influences are "alien to the agent."

Three options present themselves for addressing these facts about nonconscious processing:

1. No one is responsible for anything; control—and hence free will—is an illusion.

2. Control is all about being "reason-responsive" in a *conscious* manner. Hence control (and free will) is really independent of nonconscious processing.
3. Control in mammalian nervous systems involves both conscious and nonconscious processes. The capacity for control is seen in all mammals and probably many birds (see also Panksepp, 1998).

The first option is flatly untenable. For one thing, human communities will not cease to respond with disapproval and punishment to those who commit assault, theft, and insider trading. This is simply a pragmatic necessity of life in social groups. Holding people responsible is one element in biasing their future decisions, a nontrivial feature of all social mammals. Moreover, if you seriously try to dismantle the criminal justice system on grounds that free will is an illusion, vigilante justice will take its place. Not a desirable outcome. Dealing with miscreants is best managed within the confines of the criminal justice system. The institutional basis for assigning responsibility is deep and highly sophisticated and should continue to be relied upon. Second, parents will continue to shape their children's reward systems with approval and disapproval. Not to do so would be absurd. Biasing the child's inclinations is what social learning is largely about.

Nor does the second option pass muster. Its underlying assumption equates agency with *conscious* agency and controlled action with *consciously* controlled action. The problem with this assumption is that fails to square with the reality of human behavior and cognition. Pesky facts run up against pet philosophical ideas again. The influence of nonconscious activities on conscious decisions is demonstrable. The highly skilled actions of a basketball player, for example, happen too fast for there to be much in the way of conscious deliberation about whether to shoot or pass. Yet the player's decisions are highly controlled and voluntary. A skilled chef can see the problem with a failed cake in a flash; a skilled sea captain can see how to take the oncoming wave without going through a conscious reasoning process; a skilled obstetrician responds immediately to the realization that the infant has the cord wound around its neck. Nevertheless, all these actions display impressive control. Nor can they be explained as mere stimulus–response actions. They draw upon a vast background knowledge and competence; they draw upon skill. Adhering to the view that control is restricted to an "autonomous" domain of *conscious* reason-responsiveness and conscious deliberation is as unrealistic as a wax kettle.

These first two options are unappealing, owing to their absurd practical consequences and their failure to square with the facts of human behavior and cognition. Fortunately, however, there is a perfectly respectable sense of *control* whereby it is generally possible to distinguish between cases where a person was in normal control from those where control was compromised or absent. Bernard Madoff ran his Ponzi scheme for years, with masterful control. By contrast, the Damasio patient, E.V.R., who had extensive damage of his ventromedial prefrontal cortex, suffered compromised control following his surgery, but not before his surgery (Damasio, 1999). Control falls along a spectrum of cases, from compromised to normal, and the aforementioned two cases are separated by much distance on the spectrum. There are degrees of control and different aspects of control. The sciences of control are beginning to help us understand the mechanisms of control. They cannot give complete answers, both because the sciences themselves are manifestly incomplete but also because some issues concerning when to assign responsibility also involve judgment about social value. Thus, as we see it, the third of the options listed above has the most compelling figures of merit. Consequently it is the one we explore below.

Control and Nonconscious Processing

Questions concerning the neurobiology of conscious processes can be usefully separated into two related issues. The first concerns the mechanisms needed in order to be conscious of *anything at all*. For this function, regions of the brainstem, central thalamus, and cortex appear to play the dominant role (Llinás, 2001; Schiff, 2008; Edelman, Gally & Baars, 2011). Assuming you are conscious in this respect, then there are other mechanisms that enable you to be conscious of *something in particular*, such as seeing a face or hearing a quail call or remembering hitting a home run. If you are in deep sleep or in coma, you are not conscious of anything at all. Assuming you are awake, the particular contents of consciousness seem to be selected on a need-to-know basis, and that very selection is itself typically a nonconscious operation (Sheth et al., 2012). Not everything in your interoception or your external surroundings is in your awareness; your brain selects what is relevant, or highly salient. How this works is not well understood although progress is being made.

Conscious cognition has a rather limited capacity, and you can effectively solve only one problem at a time (*what shall I cook for dinner tonight* vs. *what shall I write for the abstract*). Nonconscious processes, by contrast, are probably not constrained by this attentional bottleneck and operate in parallel although, even at this level, some operations will be constrained to seriality by the one-action-at-a-time requirement. Conscious processes probably operate more slowly than certain nonconscious processes such as pattern recognition or associative memory retrieval. When a decision is momentous, an automatic response to the problem may be consciously reevaluated and reconsidered (see, again, Sheth et al., 2012). Even so, many day-to-day decisions are not momentous, and much decision making proceeds without much reflection. Skilled behavior, such as making bread or driving, does not usually require heavy concentration unless something goes wrong. Recognition that a decision *is* momentous or that something has gone wrong is itself typically the product of nonconscious processes of evaluation and error checking, in the way the other forms of pattern recognition are the product of nonconscious processing (Sheth et al., 2012). It is a kind of *skill*, one that can be developed as skepticism and vigilance are cultivated for certain kinds of conditions and problems. As Kahneman (2012) points out, "intuition is another name for expertise."

Furthermore, only some nonconscious processing is relevant to questions of whether or not someone has the capacity for self-control. No one, not even the most ardent defender of a reasons-responsiveness view of control, fusses about preprocessing in the retina, for example, or about maintenance of normal body temperature, both of which are managed without conscious awareness. Nobody cares that your brain nonconsciously figures out and decides where to make a saccade every 300 milliseconds, behavior that has been shown to be both goal directed and intelligent (Yarbus, 1967; Zelinsky, Rao, Hayhoe, & Ballard, 1997). And so forth for myriad other nonconscious processes, many of which involve decisions. So what subdomains of nonconscious activity *do* seem to raise a concern for self-control?

First, seemingly minor contingencies in a situation may bias us toward one action rather than another, such as helping a person who dropped her groceries or smiling at a passerby. This domain has attracted philosophers who are apt to see the nonconscious influence of situational contingencies as undermining any strong presumption of self-control even in ostensibly normal conditions (see below).

Secondly, the reward system, and reinforcement learning generally, plays a huge role in learning about the physical and social world. The reward/reinforcement system, including the basal ganglia and other subcortical and cortical structures, is crucial in the development of habits and skills whereby individuals can suppress untoward impulses, generate options, evaluate options, manage stress, rank preferences, and make decisions under temporal constraints. These operations are important in the various manifestations of control.

A third subdomain, dependent on reinforcement learning, concerns the capacity for decisions about relevance, such as the determination that a perception or a memory or an idea is or is not relevant to solving the problem at hand. Sometimes this judgment may be conscious, but even then, many irrelevant matters have already been nonconsciously excluded. Relevance determinations are selection decisions. They are related to how attention is allocated, which is often nonconsciously settled. Additionally, contexts can provide a framework for our interpretations of speech or an event, which biases us one way or another. If you are fishing at a creek and someone says, "he is too close to the bank," you take *bank* to mean *riverbank*. If you are downtown in La Jolla near the Bank of America, you will take *bank* to mean financial institution (Kahneman, 2012). Other domains may be important too, but for the purposes at hand we shall focus on these three.

Situating My Choices

In an earlier paper (Suhler & Churchland, 2009), we challenged the idea that experiments by social psychologists have shown that people normally and generally have, at best, frail control over their behavior. According to some philosophers (Doris, 2002; Harman, 1999; Appiah, 2008), human behavior is largely at the mercy of external, situational factors that cause us to respond—a position that has come to be known as *situationism*. Our actions, so goes the situationist hypothesis, are frequently mere automatic consequences of environmental factors, not the outcome of an agent with a capacity for robust control. Our habits, temperaments, character, and goals can be buffeted about by minor contingencies occurring at the moment.

As we studied the findings in question, we concluded that the situationist case is much less compelling than it initially appears once one considers the totality of scientific findings relevant to questions of control of behavior. By focusing almost exclusively on findings in social psychology, situationists have missed many other lines of research from other subfields indicating that control is nothing like as frail as situationist philosophers propose. Matters looks very different when you balance the picture with scientific data showing the robustness of control, such as the capacity to maintain a goal despite distractions, to defer gratification, to stop an action midway, to develop advantageous habits, and to suppress impulses. This is seen in humans, but also in monkeys, rats, and, one has to predict, in many other species.

The frail control hypothesis looks feeble when stacked up against the large body of findings demonstrating the pervasiveness and sophistication of nonconscious cognition and goal pursuit (Bargh et al., 1996), work on goal maintenance and its neurobiological basis (Miller & Cohen, 2001), discoveries concerning neural pathways supporting suppression of impulses of action and thought (Aron et al., 2007), and research on the limitations of conscious executive control (Baumeister, 2005). Research on "learned industriousness" is also important in demonstrating the role of the reward system in reinforcing a pattern of behavior that results in persistence in pursuing a goal (Eisenberger, Kuhlman, & Cotterell, 1992). The leading hypothesis here is that via the reward system, the sensation of high effort itself becomes rewarding. This implies not only that the capacity for robust control is real but also that it can be strengthened through reinforcement learning. This should not come as a big surprise to the situationists.

Does knowledge in a particular circumstance of possible bias, and does skill in recognizing when the situation calls for skepticism, make any difference to whether minor contingencies can exert any effect? Common sense certainly suggests so,[4] and as we discuss below, the neurobiological evidence supports common sense.

We also registered concern that in many papers cited in support by the situationists, not all subjects showed the effects, and frequently the sample size was small. We noted that by and large, the actions in questions were not hugely consequential. They concerned whether to pick up a pencil, for example, but not whether to quit law school or switch car insurers. In sum, these concerns motivate caution in drawing startling conclusions about lack of control in healthy subjects in general.

Developments in empirical psychology since the publication of our paper call for even more caution in interpreting and using the results from social psychology. There are now reports of failure to replicate results in certain of these studies.[5] Additionally, the Bakker and Wicherts (2011) study of 281 papers in psychology found that in 18% of papers, the statistical results were incorrectly reported, and in 15% the statistical results when recalculated reversed the claimed finding (Bakker & Wicherts, 2011).

Also troubling is the confession by a highly respected and widely published Dutch social psychologist, Deiderik Stapel, to concocting data for some two thirds of his publications.[6] His fraud has undermined results apparently showing that we use better table manners if a wine glass is on the table, or that people discriminate more if their environments are disordered and messy, or that we are apt to be more aggressive if we eat meat. Although these alleged results were much ballyhooed in the media, at this point there is no reason to believe they are true.

A second senior Dutch social psychologist, Dirk Smeesters, has been found guilty by Erasmus University Rotterdam of "data selection" and failing to keep adequate records after two of Smeesters's articles were withdrawn owing to analysis by Uri Simonsohn.[7] He has since resigned from Erasmus University. In his most recent article, Simonshohn (2012) describes in detail the statistical techniques used to analyze Smeester's articles, and also those by another senior social psychologist, Larry Sanna, formerly of the University of North Carolina. With great care, Simonsohn shows that in several of Sanna's articles, also, the data appear to be fabricated. At the time of this writing, Sanna has apparently resigned. One of Sanna's articles addressed the observation that there is a metaphorical relationship between morality and high ground. Sanna and colleagues suspected the relation might be literal. Therefore they tested the following prediction: Subjects higher in elevation (e.g., on a stage platform) will behave more prosocially than those on lower in elevation (e.g., in the orchestra pit), and midway between in prosociality will be those spatially located midway between the other subjects. Sanna and colleagues reported that their prediction was upheld by their data.

Even more worrying than the cases of outright statistical incompetence and data fraud are the more subtle—and likely far more common—ways in which false positives (spurious findings) can result from psychologists' choices in how they collect, analyze, and report data. These are the subject of an important paper by Simmons and colleagues (Simmons, Nelson, & Simonsohn, 2011) detailing how "researcher degrees of freedom" can dramatically increase the prevalence of false positives—the finding of an effect when in fact none exists. Critically, the degrees of freedom that they describe—such as controlling for gender, choosing one's sample size (and when to stop collecting data), and dropping certain subjects (e.g., outliers) or experimental conditions—are at times useful and appropriate tools for finding genuine effects in the data. However, when not employed with the utmost judiciousness, they can also be used to tease an effect out of the data that, at least nominally, are below the field's threshold for statistically significance ($p \leq 0.05$). Furthermore, as Simmons and colleagues' models clearly demonstrate, the probability of a spurious positive finding increases dramatically when more than one of the degrees of freedom is employed.

The aforementioned empirical difficulties do not imply that none of the results of social psychology stand. The problem, instead, is one of uncertainty. Absent third-party replication and data reanalysis, we are unsure which do and which do not stand. We do not know whether situationist philosophers relying on data from social psychologists have reanalyzed the raw data or replicated the findings or whether they just assume the conclusions are genuine. In our view, the recent spate of challenges to the data from certain social psychologists provides yet more reason to exercise caution regarding the situationist conclusions regarding frail control.

Owen Flanagan, in his careful dissection of situationism, has made some of the most devastating criticisms from a broadly philosophical and empirical perspective (Flanagan, 2009). Flanagan correctly notes that decisions are *of course* sensitive to events in the immediate environment. Is that not part of what *reason-responsiveness* involves, after all? Just as obviously, however, decisions are not overwhelmed by inconsequential events. Decisions are biased by habits, character traits, background knowledge, and temperament. In addition, decisions may be affected by other events in the agent's recent history that altered the emotional and stress valence. They may be affected by the agent's hunger or exhaustion or fear or hypothermia. As Flanagan rightly notes, facts about habits do not entail that the contingent factors must be deemed irrelevant. Such factors of the situation can affect decision making. But nor, for sure, are contingent factors *paramount* in all decision making. Decision making is a constraint satisfaction process, and habits, skills, and expertise are important constraints (Churchland, 2011; Litt, Eliasmith, & Thagard, 2008).

The Reward System: "I, Agent"

In our earlier discussion of nonconscious control, we discussed a range of empirical approaches that tend to support the hypothesis of nonconscious control. One crucial line of work we did not explore, however, was the role of the reward (positive and negative) system in establishing what Aristotle called "good habits of thought and action." Contemporary philosophers have sometimes been puzzled that Aristotle placed so much emphasis on good habits. They prefer instead to focus on conscious deliberation about "reasons" and principles in consciousness *at the time of the decision*. However, as Aristotle appreciated (and contemporary philosophers tend not to), skills—social skills, problem solving skills, decision-making skills—

are crucially important in all behavior, including social behavior, and in achieving a good life. Aristotle would have relished the recent scientific advances in understanding the reward/reinforcement system and its indispensable role in developing adaptive, socially advantageous control capacities. To this topic, we now turn.

Among many other factors, two forces have powerfully shaped brain evolution: the advantages of learning, so as to reduce uncertainty and promote adaptability, and the advantages of minimizing energy and time costs. These factors interact.

Reducing uncertainty (improving the capacity to make accurate predictions) is a fundamental constraint on the evolution of nervous systems—*all* nervous systems. Brains reduce uncertainty by learning via adjusting their circuitry to make predictions about what will happen next based on what happened earlier. They update their prediction circuitry depending on what happens now. For all animals, the capacity to make and use predictions is valuable because predictions guide subsequent behavior, aiding the animal in survival and reproduction.

To make useful predictions over all but the shortest timescales, nervous systems make structural changes to their circuitry. Growth of dendritic spines of neurons (especially those of pyramidal neurons), along with pruning back of noncontributing spines, is a major component of learning. The brain's reward/reinforcement system uses both strategies. Stable changes involving growth or pruning are seen whether the brain is learning the spatial whereabouts of about food sources or how to ride a bicycle or how to behave in a courtroom. Importantly, when there is growth, there must be gene expression to produce the proteins that constitute the wherewithal for growth. Thus learning and gene expression are tightly linked.

Learning to pursue *this* goal rather than *that* involves evaluating the expected consequences of a plan and adjusting accordingly. This must also be done in the context of homeostatic coordination of needs and drives. All mammalian brains (probably all vertebrate brains) do this. So far as mammals are concerned, the main differences across species in their reinforcement learning capacity depends on the relative complexity of the prefrontal cortex and its relationships between components of the ancient reward system (see figure 8.1). Here is an instance where quantitative difference can yield qualitative differences: Human brains have a larger prefrontal cortex, and more neurons, in absolute terms than other land mammals. This may be worth emphasizing as we aim to characterize the role of the reward system in biasing intelligent choice. By embellishing the ancient subcortical reward system organization with fancy cortical input,

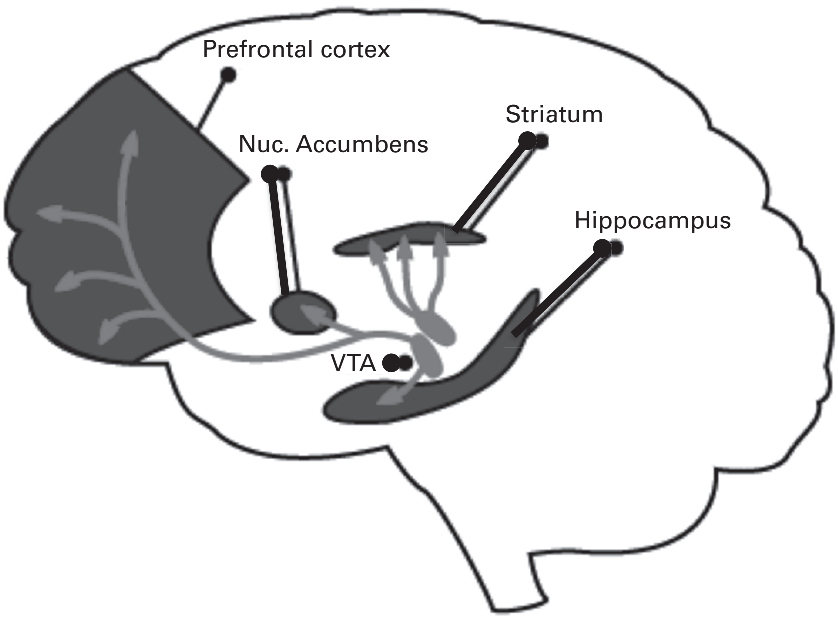

Figure 8.1
Schematic showing the main circuits of the reward system. VTA, ventral tegmental area.

a plan can be evaluated for its likely consequences. Richer cortical input allows for richer predictions and evaluations. Goals can be nested within goals. Plans can become very elaborate and goals very abstract. By drawing on learned patterns of causality, the brain can assemble evaluations of the consequences of a plan. These are "if … then" models, and they can become very sophisticated indeed.

Adjustments to circuitry reflect whether an action yielded positive or negative results. In the reward system, they also involve a prediction about what will happen on the next, similar, occasion. A positive outcome biases the brain's reward system in favor of that action; a negative outcome biases it against (see Thorndike, 1911; Hikosaka, 2010). No matter whether they have a simple or a complex nervous system, an organism's learning is contingent on the nervous system's reacting to some action outcomes as positive, to others as negative, and not reacting much at all to others. There are degrees of reward; a serious burn is really horrible, a pinprick is not so bad. If the organism's nervous system is insensitive to pain or pleasure, or if its neural circuitry cannot change (see above) even when it does care about a bad or a good outcome, it will tend to be wiped out. In our neurobiological world, the meshwork of reward and learning is a very big deal (Montague, Hyman, & Cohen, 2004).

Accurate prediction enhances the probability of survival; wasting energy does not. Animals must expend energy in finding food and water, in keeping warm and safe. Gram for gram, nervous systems are much more energy intensive than other organs. Importantly, therefore, modifying circuitry to reflect invariants in the physical and social world also helps keep energy costs down. If your nervous system learns that raspberries are tasty and choke cherries are nasty, you need not waste time and energy gathering the choke cherries. Energy savings means less time hunting randomly, less time hunting food altogether, and more time staying healthy and running after mates.

So it is in the social domain. If your nervous system is adjusted so that you avoid taking food from your older brother since he smacked you last time, you do not have to suffer pain repeatedly for stealing his food. You can find other, more energy efficient, food strategies. In humans, socialization of the young concerning what is and is not acceptable involves approval and disapproval, which means pleasure and pain. Kicking your playmate gets disapproval; including him in your soccer game gets approval. Mother Nature could conceivably have put all necessary knowledge in the genes, but it turns out to be vastly easier and vastly faster to design nervous systems to tune themselves up to the environment—to let them learn.

Our reward system is responsive to approval and disapproval, but not in a simple reflex-like way. Its response portfolio becomes ever more complex and subtle as the prefrontal cortex expands in evolution and develops during maturation of an individual. The complexity is even more daunting when you factor in coordination with homeostatic responses, stress responsivity, attachment, affiliation, and competitive risk taking. Sometimes disapproval may get downplayed if other motivations are strong enough; sometimes approval is ignored because it is perceived to be unreliable, for example. Some individuals are more stubborn than others; some are more willing to take risks. So temperamental factors play a role as well. Throughout life, but especially in the young, the slow, adaptive modification of circuitry in the prefrontal–reward pathways shapes what we feel and how we emotionally respond in social contexts (see especially Del Giudice, Ellis, & Shirtcliff, 2011). This is the development of control. Control is ancient, and in mature animals especially, it can be robust against distractions.

Classical conditioning is quite well understood, especially in neurologically simple species (Carew, Walters, & Kandel, 1981; Carew, Hawkins, & Kandel, 1983; Hawkins, Abrams, Carew, & Kandel, 1983), but until the last 15 years, the neurobiology of trial-and-error (reinforcement) learning remained elusive. Now that the basal ganglia are known to play a crucial role in this form of learning, the exact nature of that role is being extensively explored, both using single-cell recordings in animals and using imaging technology in humans (Rangel, Camerer, & Montague, 2008). The field has blossomed in a spectacular manner, though its impact on philosophy has been negligible so far.

A conceptual insight from Read Montague helps put an evolutionary perspective on why the reinforcement/reward circuitry is so important in an animal making its way in the world: In reinforcement circuitry, some of the animal's energy resources are put into responding to the reward (chowing down on the tasty fruit you found, or licking your wounds from a fight with a conspecific), but a little bit is put into modifying circuitry to embody a prediction: If next time I do A, then reward B will probably occur again. I tried, I erred, and my brain encodes that information in a bit of structure. My energy investment reaps a profit (Montague, 2006). This summary simplifies the complexity of the learning process, of course, but it illustrates the uncontroversial point that reinforcement learning was strongly selected for in the evolution of nervous systems.

Montague and his colleagues have explored an additional aspect of trial-and-error learning that goes beyond direct reward and punishment. This is *counterfactual learning*—learning based on an evaluation of the outcome of the option *not* taken (usually called *fictive error*). This is learning about what I should have done instead, and it requires a comparison between what I did get on the option taken with what I would have got on the other option. This comparison and evaluation appears to be mostly a nonconscious business, though the feeling of disappointment at the option forgone may of course be felt consciously. Its effects can be seen quite clearly in imaging studies. Notice especially that registration of fictive error is conceptually subtle, and by the any reasonable standard, *intelligent*. It is recognition of the value of the road not taken—by *me*. So in the background are the conceptual resources to relate that option to oneself, and to compare what did not happen to what did happen, as a result of *my own choice*—and then to use that comparison, perhaps but not necessarily, to adjust one's predictions and make a different choice next time. The basal ganglia and prefrontal cortex are again crucial to these computations, though at this point the precise nature of role played in this by the neurons in the prefrontal cortex remains poorly understood. By our lights, the discovery of fictive error and the possible roles its might have in how humans learn to get on in their social and physical worlds is monumental.

Whatever else agency is, it has deep roots in reinforcement learning. Surprisingly perhaps, the basal ganglia are also essential to the difficult computational task of figuring out whether it was *you* that caused a specific painful or pleasurable event and exactly what it was you *did* that produced that event (Redgrave, Vautrelle, & Reynolds, 2011). For example, the dog of PSC learned to open the latch of a gate; he pulled on the wire inside which released the unseen latch on the outside, then pushed to open the spring loaded gate. How did he know exactly which movement among the many he made was the one relevant to opening the unseen latch? You may think this information is just *given*—just obvious. It is not. From the point of view of the brain, nothing about causality is *given* in that old philosophical sense. Although it is not known exactly how the computational problem is solved, it is known that the solution depends on using precisely timed signals looped from the motor cortex to the basal ganglia that represent the motor decision you just made. Any representation of oneself as an agent that released the latch by pulling on the wire depends on computations occurring in the basal ganglia.

These functions are anchored by the sensory-driven phasic input of the neurotransmitter dopamine into the dorsal striatum, the ventral striatum (nucleus accumbens), the amygdala, and prefrontal cortex. These are the same structures that are important for the development of habits as well as social and life skills. They are important for making decisions under uncertainty, which includes pretty much everything you decide to do. They are important for learning when to slow down and reflect more carefully on what to do next. They are important for evaluating why you made a choice that was less valuable than the choice not taken. They are essential for recognizing what is and is not relevant to a problem. (See figure 8.2.)

Our hypothesis is that any animal with a healthy, functioning reinforcement learning system and a healthy system for maintaining homeostasis—a normal regime of feeling pleasure and pain, responding to environmental inputs, and modifying plans—is "reason-responsive" (again see Del Giudice, Ellis, & Shirtcliff, 2011). It is capable of generating "if … then" models of various options, even if it does not use language. A *reason*—a part of the "if … then" model—can be a perception, a memory, recognition of fictive error, a feeling, an associative connection. Reasons do not have to be cast

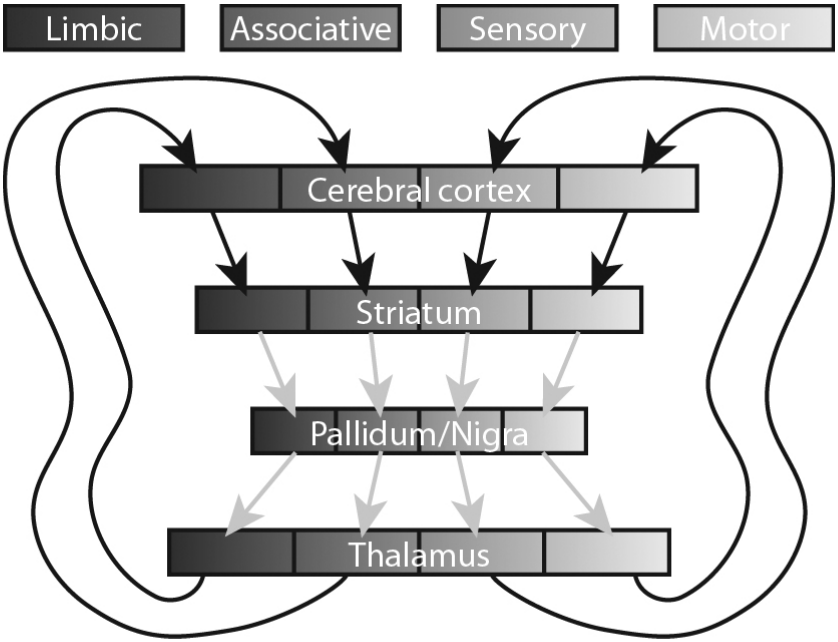

Figure 8.2
This illustrates the looping nature of the pathways between the cortex and subcortical structures and the roughly segregated nature of those pathways. The shades of gray demarcating the subcortical structures correspond to the shades of gray in the subfields of cortex. The relationship between the cortex and the striatum is very complex and still not well understood. (Courtesy of Redgrave, Vautrelle, & Reynolds, 2011; with permission.)

in a language, and they do not have to be part of a logical syllogism, though in language-using animals they could be. Reasons can get wondrously complex, as the prefrontal cortex capacity for generating models to guide behavior becomes more complex. At bottom, however, an essential function of nervous system is to coordinate all these factors in a way that is life maintaining for the individual.

In a recent paper, Kirk, Harvey, and Montague (2011) report a result that reflects back on the situationist claim while showing something remarkable about the striatum–prefrontal pathways. Roughly, the first finding is that nonexperts can be biased in favor of a piece of art when a sponsoring company's supporting logo is nearby. The second finding, however, is that experts are insulated against this bias. And not because they consciously go about insulating themselves but by virtue of their skill and past experience. Moreover, experts, but not novices, recruit an area of the prefrontal cortex—namely, the dorsolateral prefrontal cortex. Kirk et al. propose that in the experts, there is a kind of censorship by the ventromedial prefrontal cortex of the tendency to bias, the tendency seen in nonexperts. This suggests that the expert brain has evaluated the situation as conducive to bias, stimulating a kind of vigilance and skepticism that some component of the reward system probably prepares us for. This seems to us to be a wonderful example of nonconscious control, though the groundwork for it may have been conscious in early stages of learning about the world. It is also a wonderful example of how past learning can sensibly and wisely figure in how we act. It is an example of how knowledge and skill can make us more effective agents. This would not need to be emphasized save for the conviction of situationists that contingencies of the situation regularly and typically determine our choices.

The knowledge of the experts deployed in the decisions is not “alien” to them; it is part of what makes them the agents they are. Their expertise is part of who they are. It is what allows them to act intelligently in a complex situation. The knowledge does not need to be conscious when the decision is made to be effective. This point Aristotle and Hume well understood. Data from both the neurobiological and the psychological levels indicate that evaluative functions in goal generation, goal choice, action selection, assessing consequences, updating predictions, and updating goals depend on unconscious evaluative/predictive processes, *some* of whose results feed into conscious processes. These conscious processes are beholden to the midbrain dopamine system. A staggering amount of data deploying all the methods available—animal models, computer models, neurological studies of humans, imaging studies on humans, behavioral experiments—are slowly enriching an account concerning the midbrain dopamine structures as the engines of reinforcement learning. No, they are not the whole story. Yes, along with the large human prefrontal cortex, the results are hugely subtle and complex. Yes, there are other kinds of learning. And yes, there is so much we do not understand about the mechanisms. But still.

Conclusions

The evolution of the mammalian brain involves tight links between the reward (positive and negative) system and the prefrontal cortex. These links support controlled behavior, behavior that is appropriate to the life-maintaining business of the animal in its physical and social worlds.

Developing patterns of self-control during maturation is critical to the acquisition of skills that support wise choices in the social and physical worlds. Contrary to philosophers who seek a form of free will that is miraculously independent of causation, this framework is very Aristotelian. It recognizes that the realistic aim is to get the causality right—to bias goal selection and suppress impulses in a manner that serves the person well in the long haul rather than unrealistically seeking to avoid causality altogether.

But, you ask, is all that reward system activity really part of me? We reply, how can it *not* be part of you? What would you be without it? Your conscious deliberations are what they are because of how they are integrated with the rest of your brain. You are a whole, integrated individual. Your conscious life is what it is because of the way it meshes with the products of your nonconscious brain. Your habits of action and habits of thought are important for precisely the reason that Aristotle understood so well. Cultivate them carefully, make them work to your advantage, for they are a big part of what makes you *you*.

Acknowledgments

Particular thanks are owed to Peter Redgrave for his extensive and most helpful comments on an earlier draft. Many thanks are also owed to Paul Churchland and Owen Flanagan.

Notes

1. David Hume famously recognized this point in his arguments against a Cartesian conception of choice.

2. PSC made this point in *Brainwise* (Churchland 2002).

3. It is essential to avoid the common if seductive trap of trying to precisely define a concept before the science has progressed enough sensibly (not fancifully) to enable such precision. Many a good project has ground to a sad halt in the first step by the premature effort to hew out a precise definition, which condition is unfortunately is like moths to a flame for those who make a philosophical living by inventing bizarre and hilarious counterexamples to premature definitions. Precision in the definition of α and progress in the scientific understanding of α coevolve. See Churchland (2002).

4. See this discussion: http://www.sciencenews.org/view/feature/id/340408/title/The_Hot_and_Cold_of_Priming.

5. See Pashler, Coburn, and Harris (2012); Doyen, Klein, Pichon, and Cleeremans (2012). For a reply, see http://www.psychologytoday.com/blog/the-natural-unconscious/201203/nothing-in-their-heads. See further, http://www.nature.com/news/replication-studies-bad-copy-1.10634.

6. See http://www.nature.com/news/2011/111101/full/479015a.html.

7. See the report http://bps-research-digest.blogspot.ca/2012/07/has-psychologist-been-condemned-for.html. For a discussion with Simonsohn, see http://www.nature.com/news/the-data-detective-1.10937.