ResearchGate

See discussions, stats, and author profiles for this publication at: http://www.researchgate.net/publication/281452848

25 Davies 2011

DATASET · SEPTEMBER 2015

DOWNLOADS
9

VIEWS
11

1 AUTHOR:



Jaak Panksepp
Washington State University
424 PUBLICATIONS **16,477** CITATIONS

SEE PROFILE

Available from: Jaak Panksepp
Retrieved on: 04 September 2015

Neuroscience and Biobehavioral Reviews 35 (2011) 2036–2044

Contents lists available at ScienceDirect

Neuroscience and Biobehavioral Reviews

journal homepage: www.elsevier.com/locate/neubiorev

Review

Ancestral voices in the mammalian mind: Philosophical implications of Jaak Panksepp's affective neuroscience

Paul Sheldon Davies

Department of Philosophy, College of William and Mary, Williamsburg, VA 23187-8795, USA

ARTICLE INFO

Keywords:
Primary process systems
Systematicity
Executive control
SEEKING system
PANIC system
Anoetic consciousness
Causal interpretive system
Social attachment homeostasis

ABSTRACT

The philosophical implications of Jaak Panksepp's affective neuroscience comprise a significant form of skepticism regarding our capacities as agents. This is clear in two ways. (1) Panksepp's methods of inquiry support a corollary to Dobzhansky's famous maxim concerning evolution: nothing in mammalian psychology makes sense except in light of ancient affective capacities shared by all mammals. The application of this maxim, I argue, raises informed doubts concerning our knowledge of our own capacities. (2) Against the backdrop of this maxim, Panksepp's substantive discoveries provide tentative confirmation of theories in psychology which raise doubts about our alleged capacity to give reasons for our actions. Taken together, Panksepp's methods and discoveries call into question the view we have of ourselves as free and responsible agents, while pointing us toward more-fruitful forms of inquiry concerning all our animal capacities, including our capacities for deliberation, choice, and action.

© 2010 Elsevier Ltd. All rights reserved.

Contents

1. Introduction 2036
2. Human agency 2037
3. Methods for studying the human self 2037
4. Fruits of studying the human self 2040
5. Conclusion 2043
Acknowledgements 2044
References 2044

> I am inclined to agree with Francis Galton in believing that education and environment produce only a small effect on the mind of anyone, and that most of our qualities are innate.
> -Charles Darwin (in F. Darwin (Ed.), 1888/1959, p. 21)

> Our work . . . has been premised on the assumption that the evolutionary "why" for the existence of affects is that they were . . . heuristics to anticipate survival issues . . . they are the unconditioned experienced "voices" of certain genetically-constructed brain networks.
> -Jaak Panksepp (in Panksepp, 2011)

1. Introduction

Jaak Panksepp is an audacious naturalist of the mammalian mind. In the spirit of exploration that animated naturalists of the 19th century, Panksepp approaches the study of the human self with the positive expectation of unsettling current categories of knowledge. The expectation, borne not from fanciful thinking or feigned humility but from the most promising methods and theories of our day, including those developed by Panksepp, is that much of what we presently take for granted about ourselves will be undermined by progress in knowledge. Charles Darwin was the 19th century's most audacious naturalist of the history of life on earth; his discoveries changed our knowledge of all known living things. Panksepp is one of today's most audacious naturalists of the mammalian mind; his discoveries and innovations in affective neuroscience are changing our knowledge of the minds of all known living things.

The aim of this essay is to illustrate Panksepp's exploratory naturalism and describe its implications for philosophical theorizing about the human self. I defend two theses. The first is that Panksepp's methods for studying the mammalian mind, derived from knowledge in evolutionary biology, psychology, and neuroscience, force us to the conclusion that we presently have no clear

E-mail addresses: psdavi@wm.edu, psdavi@mindspring.com

0149-7634/$ – see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neubiorev.2010.10.010

idea what kinds of agents we are. Our best methods drive us toward a form of skepticism regarding human agency. I defend this claim in Section 3. My second thesis is that Panksepp's substantive discoveries force us still further toward a skeptical view of our capacities. This is true with respect to several of his discoveries, and I describe a particularly vivid case in Section 4. But I begin by extracting a few keys claims concerning our capacities as agents, which then serve as targets of the skepticism that emerges from Panksepp's affective neuroscience.

Panksepp, of course, is not a skeptic about the study of the human self. He does not recommend a general skeptical orientation of any sort when investigating the mammalian brain. His research program, moreover, is relentlessly progressive: it aspires to discover the truth about the self on the basis of accumulating knowledge of our evolutionary history, our inherited neurological capacities, and the ways our neurological inheritance implements our cognitive and social psychological capacities. All this is consistent, however, with the claim that his methods and discoveries give rise to a form of skepticism concerning human agency. As we will see, Panksepp has begun to decipher the ancestral voices in the mammalian mind, including the mind of present-day humans, and the implications of this neural archeology, at least to date, undermine substantive portions of our traditional humanistic view of ourselves. At the same time, we presently know too little to formulate a clear alternative view of ourselves – we do not yet know how to integrate knowledge of our affective capacities with theories in psychology – and therein resides the heart of our current skepticism. We know enough to turn our backs on the past but we know too little to see clearly what lies ahead.

All the same, our present skepticism is no doubt transitory. Panksepp's exploratory neural archeology, having proven its fruitfulness so far, gives us what I regard as our most powerful general framework for future studies of the human self.

2. Human agency

Which traditional humanistic assumptions are called into question by Panksepp's methods and findings? The most salient assumption is that we have, on occasion, reliable first-person knowledge of the reasons why we act. If this assumption is cast into doubt, if our first-person beliefs about our reasons for acting cannot be justified, the resulting skepticism is grave, for it calls into question associated notions of human freedom and moral responsibility. Allow me to elaborate.

Samuel Johnson famously quipped that "All experience is for the freedom of will; all theory is against it." (Boswell, 1791) This, the traditional problem of freedom of will, has been a persistently vexing conflict. The conflict has been between our experiences as agents, which tempt us to see ourselves as endowed with a powerful form of self-control and self-knowledge that renders us free, and the apparent implications of well-confirmed scientific theories, which appear to leave no room for genuine freedom of will. But, of course, the appearance of a conflict is correct only if the epistemic authority of our experiences as agents is on a par with that of our best sciences. We thus must wonder: what are the experiences which, according to Johnson, are endowed with such authority?

Such experiences surely include instances of the following general pattern: when you deliberate between two apparently possible future actions, when your deliberations accurately reflect the relative weights among your actual values and proceed without impediment, and when your reasons for one action appear to outweigh your reasons for the other, resulting in the performance of an action that matches your antecedent choice – when your deliberations work this way, you are likely to feel and see yourself as free. The match between subsequent action and prior deliberate choice is crucial. If we do not see our deliberations as causing our subsequent actions, if the reasons we give for our actions appear disconnected from our actions, then our felt sense of agency is diminished or lost. It is precisely this type of first-person experience – ones in which our actions appear to match our antecedent deliberate choices – upon which I shall focus.

Of course, such experiences constitute evidence of our freedom only if the appearance of a match is underwritten by genuine causal connections. If our first-person experiences in which our actions appear to match our prior deliberate choices are open to substantial doubt, then the alleged epistemic authority of our experiences as agents and the conflict described by Johnson are both illusory. I will argue that our first-person experiences are indeed open to such doubt. In particular, I will argue that, according to current sciences of the self, including Panksepp's affective neuroscience, it is rational to accept that

> (S) For any action we perform, we cannot justifiably claim to know, from the first-person perspective, whether our prior choices are among the actual causes of our action.

If this skepticism is defensible, then our experiences as agents have lost their former epistemic authority and Johnson's aphorism falls down dead. And, of course, if the traditional problem of free will has lost its former traction, then all philosophical theories that assert the existence of freedom of will are solutions to a problem that no longer exists. The consequences of all this for our understanding of ourselves are considerable. If all theoretical attempts to save some notion of free will are addressed to a non-existent problem, then we presently have no clear idea what kinds of agents we are.

3. Methods for studying the human self

My first thesis is that Panksepp's methods for the study of the human self are based on well-confirmed scientific theories and thus are compelling, and that they force upon us the realization that we have no clear idea what kind of agents we are. To appreciate the power of these methods, begin with the central thesis of Panksepp's affective neuroscience, namely, that the mammalian brain comprises at least seven primary process emotional systems. I discuss some of these systems in due course, but I wish to focus first on the crucial notion of a "primary process" neural system. To qualify, a candidate system must possess the following features (extracted from Panksepp, 1998, pp. 48–49):

(a) It must be developmentally entrenched.
(b) It must be functionally keyed to recurring life challenges.
(c) It must trigger a set of coordinated behaviors that addresses life challenges by activating or inhibiting sub-motor routines and autonomic-hormonal changes.
(d) It must alter the sensitivities of relevant sensory systems.
(e) Its effects must last longer than the conditions which triggered it.
(f) It must be susceptible to conditional control by emotionally neutral environmental stimuli.
(g) It must interact with cortical mechanisms implementing consciousness and decision-making.

We can summarize these features under the following three general headings.

Systematicity: A candidate system qualifies as primary only if it comprises a recurring process that operates systemically in the organism's architecture, that is, only if it comprises a set of integrated processes among distinct structures or chemical pathways. The digestive systems of mammals are "primary" in this sense, comprising the integrated effects of structures involved in perception, motor control, energy regulation, and more.

Executive control: The recurring systematic process also must be causally integrated by a mechanism which, to use an unfortunate metaphor, functions as an executive control. The notion of "executive control" should be understood more precisely in terms of a unique causal role: a mechanism wields such control only if, within a larger causal process, it functions as a potential bottleneck trigger – only if it is a relatively small or singular causal node which, when stimulated, triggers a relatively large and complex set of downstream causal processes. Some of the mechanisms comprising the digestive system are illustrative. As Panksepp notes, there are specific mechanisms in the ventromedial hypothalamus that appear to be among the "metabolic information integrators of the long-term energy balance system." They are executive in so far as the stimulation of these nodes directly triggers complicated patterns of downstream processes (Panksepp, 1998, p. 177 and the whole of chapter 9).

Recurring life challenges: The candidate system also must comprise a systemic and executively controlled process that enables the organism, often enough, to satisfy recurring life challenges. Some life challenges, such as nutrition, recur throughout the life of the organism, while others occur infrequently in the life of each organism but recur for each generation at critical developmental junctures. The distress calls of some infant mammals are illustrative. They begin when the infant becomes mobile and can wander away from its mother but they begin to taper off as the child's capacity to fend for itself begins to blossom. And such calls are the effects of a system – Panksepp calls it the PANIC system – that is indeed systematic, executive, and addressed to an early life challenge of enormous evolutionary significance.

Although we might fuss over the details of (a)–(g), we cannot seriously doubt that all mammal species comprise vital organismic processes that are systematic, executive, and life preserving. Nor can we doubt that the full repertoire of mammalian primary process systems comprises a range of vital processes that are deeply affective. Several considerations from evolutionary biology, psychology, and neuroscience support this claim. I will mention just three.

Perhaps the most general consideration comes from the historical processes known to have wrought all living things on earth. An adequate theory of any human capacity, including practical reason, must be framed in light of the fact that we are an evolved species related by descent to all extant species. The failure to factor our animal history prominently into a theory of the self is a virtual guarantee of intellectual bankruptcy. The most direct way to take account of our historical continuity with other species is to look for actual, structural continuities. Look, in particular, for homologies in that part of the organism that most fully implements its capacities for agency (or capacities that are predecessors to agency). And, having looked for homologies in the brains of countless species, it is abundantly clear that the human lower brain is homologous with lower brain structures in all reptiles and mammals, and that the human mid-brain is homologous with brain structures in all mammals.

The importance of these homologies is difficult to overstate. We know that reptiles and many mammal species today have little in the way of cortical structures. We also know that the sub-cortical structures we share with reptiles and mammals must have been sufficient for ancestral reptiles and mammals to meet their life challenges. Given this knowledge, it is plausible to suppose that these sub-cortical homologies are developmentally entrenched in present-day descendents, humans included (condition a), that they are directed toward life challenging circumstances (condition b), and that they are sufficient by themselves to trigger and coordinate life-sustaining responses to internal and external challenges (conditions c–e). We also know that the various cortical capacities of many mammal species, humans included, evolved out of these shared subcortical capacities, and this makes plausible conditions (f) and (g).

A less general consideration rests upon analogy. Consider our digestive system. Though it may strain the conceptual intuitions of some theorists, digestion is correctly conceptualized as an affective system of fundamental importance. Digestion is affective because it fulfills its systemic function, in part, by the causal efficacy of felt states within the organism. We know, for instance, that among mammal species such as rats, changes in gustatory preferences depend on felt changes produced by a loss of or return to energy homeostasis (Panksepp, 1998, pp. 181–184). Here, then, is the analogy. Since energy regulation is a significant life challenge for all mammals, and since the process of digestion comprises a set of recurring, integrated processes which control the production and consumption of energy, the digestive system is aptly classified as primary in a robust sense. By parity of reasoning, therefore, the systems that implement the more-intense affective states with which we respond to life challenging circumstances must also be classified as primary. The system that triggers separation distress calls in mammalian infants is illustrative. The PANIC system comprises a set of neural and behavioral processes controlled by a host of specific neural mechanisms (Panksepp, 1998, pp. 267–271), and it endows vulnerable organisms with the capacity to escape life threatening circumstances. A similar analogy holds between obvious primary process systems such as digestion and the other major affective systems discussed by Panksepp.

In addition, however, to neural homologies and analogies among systemic functions, perhaps the most decisive consideration is what we have learned by experimentally interfering with the brains of various species. Panksepp puts the point this way:

> The most compelling evidence for the existence of such [shared primary process affective] systems is our ability to evoke discrete emotional behaviors and states using localized electrical and chemical stimulation of the brain .... once an electrode is in the correct neuroanatomical location, essentially identical emotional tendencies can be evoked in all mammals, including humans. (Panksepp, 1998, p. 52)

We know that the anatomy and neurochemistry of our lower- and mid-brain are conserved across mammals, but we also know, on the basis of very specific electrical and chemical stimulations, that these homologous structures in us have retained the systemic functional capacities that they fulfill in other mammal species. The essentially identical emotions implemented in these homologous structures clearly include felt affective experiences. We know this because the stimulation of those circuits can serve as rewards and punishments in a variety of learning tasks. The fact that these stimulations so clearly facilitate learning is compelling evidence, moreover, that there exists in all mammals a rudimentary form of affective consciousness, an experiential feeling so primal that it is, as Vandekerckhove and Panksepp (2009) put it, anoetic or unknowing. (The existence of such elemental affective consciousness is at the heart of Panksepp's hypotheses regarding the evolution of higher forms of consciousness. See Panksepp, 1998, chapter 16; Northoff and Panksepp, 2008; Panksepp and Northoff, 2009.)

These discoveries reveal something of profound importance about the architecture of the human brain – that it comprises affective systems with inherited systemic functions shared across all mammals, that there are indeed ancestral affective voices in all mammalian minds. This is not an idle just-so story, nor a merely tantalizing theoretical expectation, but something we know from controlled experimental brain stimulations across different species. This is a point on which Panksepp rightly places great emphasis (Panksepp, 2008, 2007, 2005, 2003; the whole of Panksepp, 1998; Panksepp, 1985, 1982; etc.); the only way to effectively challenge

the existence of Panksepp's several primary process affect systems is to refute his experimental evidence on this very point.

There are of course important differences across mammals in cortical capacities, and these produce behavioral differences in the full expression of our ancient affective systems. Still, it is easy (and, alas, popular) to exaggerate the extent of those differences when theorizing about the self. Two considerations are relevant. First, if we focus on effects of brain stimulation that are more or less immediate, effects most proximate to the stimulation, we find striking similarities and few differences. We find similar immediate effects on other systems in the organism, including changes in the autonomic functions of visceral organs or changes in sensory sensitivities. Artificially stimulating the PANIC system in an infant cat or chimp will cause autonomic and sensory changes similar to those that occur in humans when the same site is stimulated. We also find similarities in whole-organism behaviors. Consider the SEEKING system. The generic function of this system is to animate the organism with a sense of anticipation or expectation, causing it to attend to and explore its environment. Yet we can stimulate this system in humans and other animals and observe the very same immediate behavior, namely, the arousal of seeking behavior (Panksepp, 1998, p. 52). This is so, by the way, even when the animal is anesthetized! (Rossi and Panksepp, 1992; Panksepp, 2011).

The second response, however, is that we can, in some instances, focus on the fully elaborated effects of the stimulation and nevertheless infer the existence of a shared affective system. To do so, we need only revert to shared systemic functions executed by homologous brain structures. Consider the stalking of cats and the sniffing of rats. If we fixate superficially on these behaviors, we will likely conclude that they are quite different. But, as Panksepp points out, cat-stalking and rat-sniffing are behavioral routines implemented in homologous brain systems and they both serve the same crucial function of locating nutrition (Panksepp, 1998, p. 194). The neural implementations are homologous and the evolutionary function served is the same; what more evidence do we need to conclude that these apparently distinct behaviors arise from a common primary process system?

These three considerations drawn from our knowledge of evolutionary biology, psychology, and neuroscience underwrite Panksepp's claim that some affective systems are robustly primary. Moreover, Panksepp's appeal to these considerations vividly illustrates the methods of inquiry with which he studies the human self. We may summarize his methods as follows, beginning with two banal biological facts:

1. *Homo sapiens* is an evolved primate related by descent to other extant primates, mammals, and so on.
2. The anatomical structures and chemical processes in the mid- and lower-human brain are homologous to mid- and lower-brain structures and processes in other mammals.

Then, on the basis of these facts, we adopt the following methodological directives when framing our inquiries concerning the human self: for any human lower- or mid-brain structure or process that is homologous across mammals, we should hypothesize the following:

I. That the homologous structure or process, if it is a primary process affective system, implements an evolved systemic function shared among mammals and dedicated to satisfying recurring life challenges,
II. A neural process qualifies as a primary process affective system if it satisfies features (a)–(g) above, and
III. The most central evolved systemic function of any primary process affective system is anticipatory in nature: primary affects prepare the organism for what is likely to occur next in its internal or external environments.

And, as I shall now explain, the adoption of these directives forces us toward the skeptical view described in the previous section. Indeed, if we take these directives as our guide, we must recalibrate many of our traditional expectations concerning the nature of our own capacities and cultivate the very different expectation that our humanistic understanding is incomplete, misleading, or, in some instances, flat-out mistaken.

To see this, consider the application of the directives in I–III. In studying the self, we begin by looking for lower- and mid-brain homologies across all mammal species that qualify as primary process affective systems, and we identify the anticipatory systemic functions of those shared systems. Once we discover neural homologies in humans, our default expectation concerning all our capacities shifts in a substantive way: we must now expect that such comparative studies give us a basic roadmap for the study of human neurology and psychology. And when we frame our inquiries in this way, as Panksepp has been doing for the last three decades, we discover the first germ of my skeptical thesis. We discover, that is, a host of substantial mis-matches between our humanistic understanding of our capacities and the actual facts.

Consider again the SEEKING system. That system constitutes us as animated, energized, and expectant explorers of the world; it is what moves us to investigate and notice things in our environment. Yet this capacity, constituted from a system that causes us to notice things, is the kind of thing that we, in our experiences as agents, are likely not to notice. Here is an analogy. Our visual experiences are implemented in mechanisms for seeing, yet we typically do not see the mechanisms with which we see. This makes good evolutionary sense, since seeing the mechanisms that enable us to see would not have done much for our ancestors' selective success. Likewise, the capacities that implement our experiences as agents are functionally dedicated to anticipating life challenging events, and it is a safe bet that first-person reflections on our capacities for anticipating life's challenges were not among the actual life challenges faced by our ancestors. So the SEEKING system is something our experiences as agents fail to reveal to us. Not just its effects, but its very existence! Nothing in our first-person experiences reveals to us the existence of this kind of system. That is a rather large mis-match between our humanistic understanding of our selves and the actual facts.

A similar point applies to all the other affective systems discovered by Panksepp. Consider again the PANIC system. It is surprising to learn that the neural system that causes us to panic in life threatening situations is also one of the main systems which move us to be the kind of social animals we are. Yet precisely that is Panksepp's hypothesis. The hypothesis is that we do not have a distinct system dedicated to social attachments in the way, for instance, our visual system is dedicated to vision. Instead, we have a system that helps free us from life-threatening circumstances that also causes us to minimize the terror of separation by constantly motivating us to construct and maintain social attachments. Here the mis-match is not about the existence of the relevant system, as it is with SEEKING, but rather the functional repertoire of the system. If the PANIC system works the way Panksepp hypothesizes, we must conclude that there is little or nothing in our experiences as agents that reveals to us the full range of anticipatory functions served by that system.

The same point can be made regarding the other systems proposed by Panksepp, including FEAR, RAGE, LUST, CARE, and PLAY: the anticipatory, systemic functions served by these primary process affective systems do not match our humanistic understanding of our own capacities, at least along some very substantial dimensions. We thus may generalize and conclude that our experiences as agents are not a reliable source of knowledge about our primary

process systems. But – that is only part of the story, because it turns out that our agential experiences are also a poor source of knowledge concerning our non-primary capacities. To see this, consider two additional methodological directives:

IV. The study of our non-primary process affective capacities must be framed in terms of knowledge of our primary process affective systems, and
V. The study of our non-affective capacities (if any) must be framed in terms of knowledge of our primary process affective systems.

IV is plausible in light of our evolutionary history and psychology. We know that our non-primary affective capacities evolved out of and depend causally on our primary process affective systems. That is why some non-primary emotions appear so similar to some that are clearly primary. Shame, for example, is sometimes highly visceral, difficult to modulate or extinguish, and vital to social attachments. At the very least, shame appears to be constituted quite directly from central features of some primary emotions. It appears, for example, to be the product of the felt threat of separation distress implemented in the PANIC system, triggered by the representation of oneself as failing to meet some social norm.

The directive in V is ambiguous. It appears to presuppose that some capacities of the mind are entirely non-affective, and that seems unlikely. A more plausible reading is that some capacities of the mind simply depend less for their operations on affective capacities than others (Liotti and Panksepp, 2004; Northoff et al., 2004). On either reading, however, we can be confident that all our cognitive capacities depend for their operations on the causal efficacy of primary affective systems. We know, for instance, that life-sustaining affective processes can occur even in the absence of cortically-based processes. Animals that have had their cortex removed early in life exhibit surprisingly coherent behavior, coherent enough to fool college students into believing that they (the decorticated animals) are neurologically intact (Panksepp, 2008, 1998; Panksepp et al., 1994; etc.). The same is true of humans whose cortical structures fail to develop: despite enormous deficits in learning and cognition, they exhibit behavior that is affectively coherent (Shewmon et al., 1999). Moreover, with respect to our capacities for practical reason in particular, the causal dependence of our actions and choices on our affective capacities is well-supported by the work of Damasio and his colleagues (Damasio, 1994).

This, then, is where we stand. Our experiences as agents are not a reliable source of knowledge concerning our primary process systems; this is something we know in light of the directives in I–III and the resulting mis-matches we discover. Moreover, since our non-primary capacities are causally dependent on the workings of primary systems, our experiences as agents are also an unreliable source of knowledge concerning of our non-primary capacities. This is something we know in light of the directives in IV and V. We know, in particular, that it is a methodological mistake to theorize about any psychological capacity from a perspective that is mainly top-down. That is a mistake because, in the absence of empirically-informed hypotheses concerning the affective systems entrenched in the mammalian brain, we are likely to theorize about phenomena that do not exist or, at minimum, phenomena about which we should be skeptical. (This applies to all our alleged mental capacities, including the higher forms of consciousness that tend to serve as the main "phenomena" about which many philosophers theorize.) Our experiences as agents, in sum, are an unreliable source of knowledge for most or all the capacities in our psychological repertoire, including the capacities that inspired Johnson's aphorism.

This is not to deny the apparent truth of several apparent facts about ourselves. We do indeed have experiences that feel from the first-person point of view to be highly cognitive and bereft of affect. Solving a math problem can feel that way. (This sort of feeling is unreliable all the same. Experiments show that arousal in the SEEKING system is affectively valenced even when we do not readily experience it as so. See Panksepp, 1998, p. 149.) Nor do I deny that we have experiences of non-primary process emotions such as shame or jealousy which would not exist absent key cortical processes that render us responsive to social expectations. The skeptical claim, rather, is that we cannot justifiably claim to know, on the basis of our experiences as agents, the correct nature of the many experiences implemented by our primary and non-primary capacities. We think and feel that we know what our various experiences are like and, of course, some of what we think and feel is surely correct. The problem is that at least some of the substantive things we think or feel about our capacities are demonstrably false – that is the upshot of applying Panksepp's methods of inquiry – and, since we cannot discriminate on the basis of our first-person experiences between thoughts and feelings that are correct and those that are not, the reliability of our agential experiences as a whole is called into question.

4. Fruits of studying the human self

The second thesis I wish to defend builds upon a central lesson from the previous section: to discover the truth about any psychological capacity, including our less-affective capacities, we must frame our inquiries in terms of what we are learning about our primary affective capacities. Dobzhansky (1973) rightly said that nothing in biology makes sense except in terms of evolution, and an important corollary is that nothing in mammalian psychology makes sense except in terms of shared primary process affective systems. My second thesis, then, is that some of Panksepp's substantive discoveries in affective neuroscience force us still further toward a skeptical view of our capacities as agents. They do this by integrating with skeptical theories from psychology – theories that call into doubt our alleged capacity to give reasons for our actions – and providing tentative confirmation in the form of plausible affective mechanisms. I will show this by focusing on the joint effects of the SEEKING and PANIC systems.

I wish to emphasize that the confirmation involved is indeed tentative. The distance between current knowledge in affective neuroscience and in relevant psychological theories is simply too great; we do not yet have sufficient theoretical connective tissue to fill the gap. That said, some confirmatory frameworks are more plausible than others. This is where the corollary to Dobzhansky has some force. If nothing in mammalian psychology makes sense except in terms of shared primary process affective capacities, then our alleged capacity to give reasons for our actions is no exception. Unless our primary process affective systems are regarded as foundational to the evolved constitution and operations of our alleged reason-giving capacity, then little of what we currently assume about that capacity will survive the accumulation of further knowledge. And that is how Panksepp's substantive discoveries should be understood, as providing a powerful affective lens through which our alleged reason-giving capacity can be seen, even if tentatively, as the effects of various primary process affective capacities, including effects that can lead us to see ourselves as something we are not.

I begin, then, with the generic assumption that the human mind, like all mammalian minds, comprises some mechanism or other for causal intelligibility, a mechanism that enables an organism to formulate internal anticipatory states concerning causal relations. If this were not the case, if organisms were not disposed to expect a certain kind of event upon observing a prior kind of event, we would be left to wonder how anyone manages to navigate the world for nutrition, safety, and the rest. A slightly less generic assumption is that the human mind, like the minds of all social mammals, comprises a mechanism that generates causal intelligibility specifically

among its social relations. Such a mechanism enables social organisms to formulate internal anticipatory states concerning relations between themselves and other organisms. Again, if this were not the case, if the causal intelligibility required for social interactions were not a high priority in the economy of our psychology, the attachments with which we live would be difficult to explain.

The precise forms of causal intelligibility depend on the range of organismic capacities. Among several primate species, including our own, social capacities apparently include expectations regarding various forms of reciprocity and perhaps even a sense of fairness. Reciprocity is evident in adult chimpanzees who are more likely to share food with other chimps who groomed them earlier in the day (De Waal, 1989). And a sense of fairness appears in capuchin monkeys who withdraw from a social exchange when they observe others receiving a better reward for the very same task (De Waal, 1996). In both cases, the causal intelligibility involved surely includes internal anticipatory states concerning, for instances, the kinds of behaviors or rewards expected in exchange for prior favors or exchanges. This is clearly true of us. If our employer pays me more than it pays you for the same service, you will find it hard not to wonder why. You will also find it hard not to resent me and certain persons in management.

Of course, we humans appear unique in our ability to address issues of reciprocity and fairness by way of arguments; we appear unique in our capacity to give and be moved by reasons. Be that as it may, we may observe that, just as certain expressive mechanisms in the capuchin reveal its frustration over an unequal distribution of rewards, a wider range of expressive mechanisms accomplish much the same in us. One very prominent mechanism is verbal articulation. Your abundant capacities for verbal behavior are likely to call me or someone in management to account for the inequity in compensation; you are likely to demand the reasons why such an arrangement is tolerated. We exercise our expressive mechanisms and impose them upon one other in virtually all our interactions. It is also something we expect mature humans to impose on themselves; an adult is judged to be immature if she fails to hold herself to account for actions she performs.

Here, then, is where the psychological theories of interest get their traction. What are the psychological mechanisms that implement our capacity to call ourselves and others to account for what we do? What, more specifically, are the mechanisms that implement our capacity to identify the causes of our actions? You and I would be unable to maintain relations of reciprocity and fairness if we did not have the capacity to identify reasons for our actions. And since a "reason" for an action, whatever else it might be, must be a causal factor for the action, the question is a pointed one: how in fact do we identify the causes of our actions?

According to the theory developed by Wilson (2002) and Nisbett and Wilson (1977), the mechanisms with which we identify the causes of our actions are explanatory: they work by generating an inference to the best explanation based on information concerning the action and the situation in which it is performed. (This is but one theory among a large number of theories which collectively cast doubt on our alleged reason-giving capacity. Wegner (2002) unifies several such theories with his own doubts concerning our capacity to consciously will our actions.) Suppose you are accosted on the street by someone begging for money, and you give her five dollars. After recounting this episode to me, I ask you why you acted as you did; I call on you to account for your action. According to Wilson's theory, my question triggers in you a causal interpretive system dedicated to rendering your actions intelligible. This system is triggered when it receives

1. The perception of yourself performing an action in a specific situation.

Armed with this information, the interpretive system begins to execute it core function:

2. The system searches your memory for causal generalizations that appear relevant to the action and situation.

Relevance depends on details. The causal generalizations retrieved from memory will likely depend on whether the person begging for money appeared sober, whether her plea for help rang true, whether you happened to have sufficiently small denominations, and so on. And the operations of your retrieval mechanism will no doubt depend on your affective states. If you are frightened or if you become angry, the generalization retrieved may be relatively blunt, but if you are calm and have time to weigh details of the situation, the generalization retrieved by may be relatively nuanced. Having retrieved a relevant generalization, the system then completes its function:

3. The system applies the generalization (step 2) to the action and situation (step 1) and draws an inference concerning the cause of the action.

These three steps may occur with great speed at a level of processing beneath conscious awareness. When, however, your interpretive system is triggered by my question, the conclusion drawn in step 3 will likely rise to conscious awareness. When that happens, when the cause cited enters conscious awareness, you have the experience of having identified the "reason why" you acted as you did. ("Conscious awareness" refers to whatever forms of consciousness required for the exercise of our alleged reason-giving capacity. It is more complex than the anoetic affective consciousness mentioned above, for it plausibly involves self-consciousness. Yet, in light of the previous section, it is plausibly implemented, in part, in the mechanisms that constitute our anoetic affective consciousness.)

Much could be said in defense of Wilson's interpretive system – I review three general categories of evidence in support of some such system in Davies (2009, chapter 7) – but I wish to emphasize the following point. The operations of this system typically occur in situations involving real or imagined social relations; the "why?" questions which trigger the system are the voices of other agents calling us to account for what we have done. But the social expectations that affect our behaviors are not always simple or consistent and, worse yet, they are often inaccessible to conscious awareness. And this is a source of potential mischief. Consider an action in which you have little time for conscious deliberation, when an immediate decision is forced upon you. Your interaction with the begging woman may be like this. Yet, when I ask you why you helped, the operation of your interpretive system is no doubt affected by what you think of and feel for me (or anyone else who is part of our conversation). At the time of your action, your thoughts and feelings for me were probably irrelevant, but they become relevant when I later call on you to account for your action. More generally, because our interpretive systems are triggered when we are called to account for what we have done – typically after the action is completed – the operations of our system are potentially affected by a host of non-conscious social factors that are distinct from factors operating at the time of the relevant action.

Still other forms of error are possible. Suppose you receive a request in the mail to donate money to a local homeless shelter and, after much conscious deliberation, you make a generous donation. When I later ask you why you decided to make a donation, your interpretive system produces a causal inference in the usual way. But suppose that, in this case, your interpretive system also came online prior to your conscious decision. Suppose, that is, that you learned to trigger your interpretive system simply by imagining

yourself performing a specific action, and suppose you learned to use the imaginative outputs of your system to uncover the reasons you might have for a given action. This would be to adapt a psychological capacity that typically operates retrospectively in the service of deliberating about future actions. All the same, the potential for error is clear. The source in this case is not a temporal gap in the effects of our affective capacities, as in the previous case, but a gap between effects that rise to conscious awareness and those that do not. As we saw in the previous section, at least some effects of our primary affective systems tend not to appear in our conscious deliberative fields; some affective effects are no doubt inaccessible and others are just difficult to accurately gauge. Either way, the output of your interpretive system appears as a "reason why" only when it enters conscious awareness. And that means that the effects of our affective capacities that do not or cannot rise to conscious awareness are effects that your interpretive system is bound to miss. This is a source of potentially enormous error in the reasons we give for our actions.

These (and other) forms of error provide considerable support for the following epistemic hypothesis: for any action we perform, we cannot justifiably claim to know, from the first-person perspective, whether our prior choices actually caused our action. This, of course, is the skeptical thesis (S) introduced earlier. And notice that the thesis is not the product of merely logically possible forms of error, but of errors that may readily occur within the operations of a system affected by various factors. Notice too that the worry does not assume that our causal interpretive system always or even regularly leads us astray. It assumes only that the system is indeed prone to error and that we are unable to discriminate, from within our first-person experiences as agents, between cases in which our system errs and cases in which it does not. But the most important thing about this skeptical worry is that the potential for error arises from the effects of our affective capacities. This is a crucial point that tends to go unremarked by psychologists of the self, including Wilson. Wilson appeals mainly to the misleading effects of mechanisms posited in the heuristics and biases literature, mechanisms that are largely cognitive in their operations. But, as we saw in the previous section, a mainly cognitive orientation for studying the self must be framed with knowledge of our primary process affective systems. I am not suggesting that Wilson is blind to the effects of affective capacities; that is surely not so. The point is one of theoretical orientation. For Wilson, affects can affect the operations of core psychological systems, but the core systems themselves consist mainly of non-affective, cognitive mechanisms. Just consider the three-step interpretive system described above. By contrast, for Panksepp, affects are foundational to all core psychological systems – consider again the corollary to Dobzhansky's maxim – and hence must be among the central mechanisms posited in any adequate model of even the most clearly cognitive system.

Consider, then, the two most-general systemic functions of the PANIC system. The first is that it enables the organism to extract itself from threats to its life. For most mammalian infants, the greatest threats to its life occur when it is separated from its mother; hence the widespread production of separation distress calls. Beyond infancy the greatest threats among mammals include predation, infanticide, extreme pain or hunger, choking, drowning, and so on – situations in which a panic response increases the chances of survival. The second systemic function of PANIC is a critical instance of the first, for one of the most significant threats that many mammals face is the loss of significant social attachments. Extreme loneliness or grief can cause us to panic as much as extreme pain or predation. What this suggests, and what Panksepp hypothesizes, is that the second function served by the PANIC system is to motivate the organism to engage in and maintain social attachments that bear on survival and reproductive success: "The most reasonable supposition at present is that social bonding ultimately involves the ability of young organisms to experience separation distress when isolated from social support systems and to experience neurochemically mediated comfort when social contacts are reestablished" (Panksepp, 1998, p. 274).

This is to conceptualize PANIC as a neural system that produces intensely negative feelings when our social attachments are threatened and that produces positive feelings, some very intense, when those attachments are reestablished or reinforced. The systemic function of these intense affective states and thus of the PANIC system as a whole is best conceptualized as the maintenance of a *social attachment homeostasis* essential for health and life. This is to conceptualize the mammalian brain as implementing a set of affective systems that move it intrinsically to form a homeostatic balance with other organisms similarly structured in their affective capacities. And it is to conceptualize the whole animal, certainly the whole human animal, as neurologically and psychologically incomplete without substantive emotional ties to significant others.

It is within the theoretical framework of the PANIC system that we may return to the earlier discussion of reciprocity and fairness, and to the requisite capacity for understanding our own actions in causal terms. A causal interpretive system like the one proposed by Wilson appears indispensable for social animals like us. Causal intelligibility enables us to account for our actions when called to do so, which enables us to engage in and maintain relations that are reciprocal and fair, which contributes substantially to fulfilling the systemic function of the PANIC system. To accomplish this, however, the PANIC system needs help. It needs help from the mechanisms with which we seek causal relations in our social relations. And that brings us back to the SEEKING system. That system responds unconditionally to "homeostatic imbalances (i.e., bodily need states) and environmental incentives" (Panksepp, 1998, p. 145) and probably participates in every emotional process (Burgdorf et al., 2007). It achieves these general functions by animating the organism to re-establish and maintain various forms of internal homeostasis. Crucially, when the form of homeostasis at issue is the maintenance of social attachments, the SEEKING system animates the organism to escape social isolation by establishing social contact and, in good anticipatory fashion, it animates the organism to prevent social isolation by feeding one's social attachments. And one way to feed one's social bonds, at least among humans, is to be vigilant about the forms of reciprocity and fairness upon which many of our relationships depend. Vigilance, in turn, requires a host of psychological capacities, including the capacity to give reasons for one's actions when called to account. And that, as we have seen, requires the capacity to identify the causes of one's own actions.

Here, then, is the second germ of skepticism. As Panksepp describes it, the SEEKING system moves us to engage in anticipatory causal learning which sometimes leads us to expect causal relations even where none exists! Causal delusions sometimes arise from the normal functioning of the SEEKING system in the service of feeding our social attachments. To see the potential for causal delusions, consider a simple case of non-social learning. Suppose we train an animal to expect a squirt of tasty fruit juice 1 s after seeing a flash of light at 20 s intervals. Suppose that early in its training the animal's dopamine level spikes upon receipt of the tasty treat and remains high throughout the 20 s intervals. Suppose further that, as training progresses, the dopamine levels drop off quickly after receipt of juice and then creep back up prior to, in anticipation of, the next delivery of juice. Well, Wolfram Schultz trained chimpanzees in just this way and discovered that dopamine levels do indeed spike immediately after the flash of light and immediately before the delivery of juice (Schultz, 2001). The chimps, it seems, learned to anticipate juice as an effect, perhaps a causal effect, of the flash of light. That the chimps had learned to anticipate juice

as a causal effect of the illuminated light is confirmed by a similar experiment conducted with pigeons. After learning to expect food immediately after a flash of light, some pigeons apparently decided to take matters into their own hands (beaks, really) by pecking vigorously at the illuminated light source. They began to behave as if they believed they could causally trigger the delivery of food even though, as a matter of fact, the flash of light did not cause the delivery of food and their pecking did not cause the light to illuminate (Williams and Williams, 1969).

The pigeons' behavior, of course, is delusional and such delusions are neither rare nor marginal; they are byproducts of a core system observed in all mammalian species that have been studied (Panksepp, 1998, p. 161). The system that moves us to seek connections between events that bear on our well-being also moves us, under certain conditions, to expect connections that do not exist. That the system does this for hungry pigeons is perhaps easy to fathom, since acting on some causal expectations or other, even false ones, may increase the relative probability of avoiding starvation. And the same might be true for other forms of homeostatic imbalance. In particular, when the social attachments of a human being are under threat, it might be the better part of evolutionary wisdom for the organism to posit causal expectations, even false ones, if such expectations diminish the chances of isolation and death. Indeed, if ancestral mechanisms operating in these ways were selectively advantageous, then mechanisms that form false causal beliefs in social situations may be built into the architecture of our minds. They may serve their social systemic functions only when the actual causal facts tend not to rise to conscious awareness. The drive to maintain social attachment homeostasis may be a fertile source of delusions concerning our capacities as agents.

This, of course, is a large-canvass speculation. We know that the SEEKING system moves us under certain conditions to acquire causal expectations that are false, but we do not know the precise range of these conditions or whether they obtain in the case of Wilson's interpretive system. However, the following circumstantial evidence, I believe, gives us our best theoretical bet at present. (1) We have behavioral evidence in support of Wilson's interpretive system; it is rational to accept that such a system exists because its existence best explains the wealth of psychological experiments (reviewed in Davies, 2009, chapter 7). (2) The effects of the SEEKING system, when contributing to homeostasis in energy regulation, demonstrably produce causal delusions, and that is prima facie evidence that the SEEKING system is capable of producing similar delusions when contributing to homeostasis in the operations of the PANIC system. (3) Therefore, one plausible neural-level explanation of the errors to which Wilson's system is prone is that the SEEKING system, in subordination to the PANIC system, is calibrated to error on the side of causal delusions when searching for ways to reestablish or reinforce homeostasis in our social attachments.

If this speculation is on track, then the psychological experiments that appear to confirm Wilson's system can be tested further: if his behavioral experiments cohere well with what we know about our primary process affective capacities, then the existence of his interpretive system is further confirmed. And the considerations discussed in this section suggest that such coherence does indeed obtain. If so, then the skepticism concerning our alleged capacity for giving reasons for our actions must be taken seriously. Our first-person experiences as agents do not have the epistemic authority that Johnson and most of our intellectual predecessors have assumed; our experiences as agents, that is, do not constitute reliable evidence for any kind of freedom of will. The conflict in Johnson's aphorism has been killed by the growth of knowledge. The conclusion to draw from Panksepp's substantive discoveries, when integrated with theories in psychology, is that the intellectual problems we face today in understanding the human self are significantly different from those faced by our Enlightenment predecessors. All theory is against freedom of will and, as we frame our cognitive theories in light of knowledge of our primary affective systems, fewer and fewer of our experiences as agents tell in favor of freedom of will.

5. Conclusion

In these two ways, then, Panksepp's exploratory naturalism forces us toward a substantive form of skepticism concerning our capacities as agents. The first is an effect of the background knowledge with which Panksepp frames his inquiries. Knowledge from evolutionary biology, psychology, and neuroscience supports the methods summarized in Section 3, and those methods, applied to the study of the human self, lead to the discovery of important mis-matches between our traditional humanistic views and the actual facts. These mis-matches, I believe, are sufficient to warrant the skeptical thesis expressed in (S). Indeed, once we grant our genealogical relatedness to other mammals, the power of Panksepp's methods becomes obvious, and once we apply those methods to the study of ourselves, our prior confidence that we know our most important capacities as agents is badly shaken.

The second route from Panksepp's work to my skepticism is far more speculative than the first, but it is potentially more devastating. The speculative nature of this route is unavoidable given our current limitations: we simply know too little to integrate psychological theories of our conscious capacities directly into Panksepp's affective neuroscience. Still, in light of Panksepp's substantive discoveries, it is a safe bet that some theory or other concerning our primary process affective capacities must be a central part of the full story concerning our capacities as agents. It is a good bet, in particular, that PANIC and SEEKING will be core theoretical posits in any adequate theory of our capacities for practical reason. If so – if the capacities with which we give reasons for our actions include the effects of ancient affective systems – then the second, more-virulent line of skepticism that emerges from Panksepp's work is difficult to reasonably deny.

We have, then, a clear methodological choice for the study of ourselves. The most venerable approach is exemplified in Johnson's appeal to the way things appear to us from our first-person experiences as agents. Experiences in which our actions appear to match our prior choices are taken as authoritative – so authoritative, in fact, they are said to match the authority of our emerging scientific worldview and give rise to the traditional problem of free will. This approach, alas, is an orientation towards inquiry that is conceptually conservative or conceptually imperialistic, or both, and that continues to retard progress in knowledge of the self. (See Davies, 2009.) The second approach is Panksepp's neural archeology. We begin not with our experiences as agents but rather by incorporating Darwin's historical insights into our general framework: we search for homologies in that part of the body most central to our capacities as agents. And once we discover homologous neural systems that implement our primary affective capacities, we begin to decipher the ancestral voices in the mammalian mind. We thus discover, as audacious naturalists of the mammalian mind, that we humans are unwitting hostages to those things that rise to conscious awareness, that we unwittingly fail to credit much that does not come to conscious awareness, and that we are particularly oblivious to the effects that our ancestral affective capacities exert upon our less-ancestral and less-affective capacities as practical agents.

We discover, in short, that the growth of knowledge in neural archeology has defeated Johnson's approach and that, at present, we have no clear idea what kinds of agents we are.

Acknowledgements

I am grateful to Vern Bingman for inviting me to participate in the Jaak Panksepp Festschrift Conference at Bowling Green State University, May 21st–22nd, 2010, where I presented an ancestral version of this essay. Thanks to efforts of Vern and his co-organizer, Casey Cromwell, and to the many Panksepp colleagues and students who presented synopses of their current research, the conference was enjoyable and exceptionally stimulating. I am grateful to Jaak for helpful comments on the penultimate draft of this essay, and to two anonymous referees for their helpful comments. My title was inspired by Ross Buck's phrase "ancestral voices of the genes" (Buck, 1999) which Jaak quotes with approval in several of his writings.

References

Boswell, J. 1791 Life of Johnson. Oxford University Press, London (1970).
Buck, R., 1999. The biological affects: a typology. Psychological Review 106 (2), 301–336.
Burgdorf, J., Wood, P., Kroes, R., Moskal, J., Panksepp, J., 2007. Neurobiology of 50-kHz ultrasonic vocalizations in rats: electrode mapping, lesion, and pharmacology studies. Behavioral Brain Research 182, 274–283.
Damasio, A., 1994. Descartes' Error: Emotion, Reason, and the Human Brain. G.P. Putnam's Sons, New York.
. Darwin, F. (Ed.), 1888/1959. The Life and Letters of Charles Darwin, Including an Autobiographical Chapter, vol. I. Basic Books, Inc, New York.
Davies, P., 2009. Subjects of the World: Darwin's Rhetoric and the Study of Agency in Nature. University of Chicago Press, Chicago.
De Waal, F., 1989. Food sharing and reciprocal obligations among chimpanzees. Journal of Human Evolution 18, 433–459.
De Waal, F., 1996. Good Natured: The Origins of Right and Wrong and Humans and Other Animals. Harvard University Press, Cambridge, MA.
Dobzhansky, T., 1973. Nothing in biology makes sense except in light of evolution. The American Biology Teacher 35, 125–129.
Liotti, M., Panksepp, J., 2004. On the neural nature of human emotions and implications for biological psychiatry. In: Panksepp, Jaak (Ed.), Textbook of Biological Psychiatry. Wiley, New York, pp. 33–74.
Nisbett, R., Wilson, T., 1977. Telling more than we can know: verbal reports on mental process. Psychological Review 84, 231–259.
Northoff, G., Heinzel, A., Bermpohl, F., Niese, R., Pfennig, A., Pascual-Leone, A., et al., 2004. Reciprocal modulation and attenuation in the prefrontal cortex: an fMRI study on emotional-cognitive interaction. Human Brain Mapping 21, 202–212.
Northoff, G., Panksepp, J., 2008. The trans-species concept of self and the subcortical-cortical midline system. Trends in Cognitive Sciences 12, 259–264.
Panksepp, J., 1982. Toward a general psychobiological theory of emotions. Behavioral and Brain Sciences 5, 407–467.
Panksepp, J., 1985. Mood changes. In: Vinken, P.J., Bruyn, G.W., Klawans, H.L. (Eds.), Handbook of Clinical Neurology (Rev. series). vol. 1. (45): Clinical Neuropsychology. Elsevier Science Publishers, Amsterdam, pp. 271–285.
Panksepp, J, 1998. Affective Neuroscience: The Foundations of Human and Animal Emotions. Oxford University Press, New York.
Panksepp, J., 2003. At the interface of the affective, behavioral, and cognitive neurosciences: decoding the emotional feelings of the brain. Brain and Cognition 52, 4–14.
Panksepp, J., 2005. Affective consciousness: core emotional feelings in animals and humans. Consciousness and Cognition 14, 30–80.
Panksepp, J., 2007. Emotional feelings originate below the neocortex: towards a neurobiology of the soul. Behavioral and Brain Research 30, 101–103.
Panksepp, J., 2008. Carving 'natural' emotions: 'kindly' from bottom-up but not top-down. Journal of Theoretical and Philosophical Psychology 28, 401–428.
Panksepp, J., 2011. In defense of multiple core affects. In: Zachar, P., Ellis, R.D. (Eds.), Dimensional Versus Categorical Models of Affect. John Benjamins, Amsterdam.
Panksepp, J., Normansell, L.A., Cox, J.K.F., Siviy, S., 1994. Effects of neonatal decortication on the social play of juvenile rats. Physiology & Behavior 56, 429–443.
Panksepp, J., Northoff, G., 2009. The trans-species core SELF: the emergence of active cultural and neuro-ecological agents through self-related processing within subcortical-cortical midline networks. Consciousness and Cognition 18, 193–215.
Rossi III, J., Panksepp, J., 1992. Analysis of the relationships between self-stimulation sniffing and brain-stimulation sniffing. Physiology & Behavior 51, 805–813.
Schultz, W., 2001. Reward signaling by dopamine neurons. Neuroscientist 7 (4), 293–302.
Shewmon, D.A., Holmes, G.L., Byrne, P.A., 1999. Consciousness in congenitally decorticate children: developmental vegetative state as self-fulfilling prophecy. Developmental Medicine & Child Neurology 41, 364–374.
Vandekerckhove, M., Panksepp, J., 2009. The flow of anoetic to noetic and autonoetic consciousness: A vision of unknowing (anoetic) and knowing (noetic) consciousness in the remembrance of things past and imagined futures. Consciousness and Cognition 18, 1018–1028.
Wegner, D., 2002. The Illusion of Conscious Will. The MIT Press, Cambridge, MA.
Williams, D., Williams, H., 1969. Auto-maintenance in the pigeon: sustained pecking despite contingent non-reinforcement. Journal of the Experimental Analysis of Behavior 12, 511–520.
Wilson, T., 2002. Strangers to Ourselves: Discovering the Adaptive Unconscious. Harvard University Press, Cambridge, MA.