

Knowing How it Feels: On the Relevance of Epistemic Access for the Explanation of Phenomenal Consciousness

Itay Shani

Kyung Hee University

Consciousness ties together knowledge and feeling, or sapience and sentience. The connection between these two constitutive aspects — the informational and the phenomenal — is deep, but how are we to make sense of it? One influential approach maintains that sentience ultimately reduces to sapience, namely, that phenomenal consciousness is a function of representational relations between mental states which, barring these relations, would not, and could not, be conscious. In this paper I take issue with this line of thought, arguing that neither of these salient aspects of consciousness reduces to the other. Instead, I offer an explanatory framework which takes both sentience and sapience as ontological fundamentals and explore how they co-evolve. In particular, I argue that while epistemic access cannot generate experience from scratch it does play a crucial role in constituting an important form of higher-order experience, namely, the capacity to experience a sense of ownership over one's experiential domain.

Keywords: sapience, sentience, panpsychism

“Herein lies the great mistake of the Cartesians, that they took no account of perceptions which are not apperceived.”

Leibniz
Monadology

Consciousness is a multifaceted phenomenon, and the concept of consciousness is a mongrel connoting a variety of different senses (see Block, 1995; van Gulick,

This work was supported by the *National Research Foundation of Korea Grant funded by the Korean Government* (NRF 2010–220–A0001). A much earlier and considerably different version of the paper was presented on February 9, 2013 at the *Consciousness and Intentionality: Franz Brentano's Heritage in the Philosophy of Mind* conference in Salzburg, Austria. I thank the participants and organizers of this conference and in particular Johannes Brandl, Uriah Kriegel, and Elisabetta Sacchi, for their thoughtful comments. I also thank Liam Dempsey, an unknown reviewer, and the editor of this journal, Raymond Russ, for their wise advice. Correspondence concerning this article should be addressed to Itay Shani, Ph.D., Department of Philosophy, Kyung Hee University, 1 Hoegi dong, Dongdaemun gu, Seoul, 130–701, Korea. Email: ishani479@hotmail.com

2014). It is widely assumed that the puzzle of consciousness is the puzzle of fitting this multifaceted and remarkable phenomenon into the wider nexus of reality, finding it a place in nature, as it were (see, for example, Rosenberg, 2004). Yet, it is not hard to see that part of the puzzle is also internal, namely, that the challenge consists, in part, in the difficulty of bringing the various facets of consciousness into mutual accord. When the task emphasized is that of integrating consciousness with the rest of nature we ask questions such as: Can the raw feels of conscious experience be nothing but physical processes? Or, can the subjective dimension of consciousness be a part of an objective physical order? In contrast, when emphasis is laid on the internal task of bringing the various facets of consciousness into mutual accord the relevant query is of a different type, to wit: What has one aspect of conscious experience, X, to do with another aspect Y? It is this latter sort of problem which occupies me here. More specifically, my goal is to investigate the nature of the connection between two of the most fundamental features of consciousness: the phenomenal dimension of felt experience, and the cognitive dimension of knowledge and information processing.

Feeling and knowing are, without a doubt, among the most recognizable, general, and fundamental features of consciousness. When conscious, we *experience*, and there is a felt quality, a phenomenal character or “something it is like,” to our experience. But, when conscious, we are also *aware* of something, which is to say that the act of experiencing is also an act of knowing, laden with epistemic qualities. Philosophers refer to these two aspects, or features, as “sentience” and “awareness” (de Quincey, 2002), “experience” and “information” (Flanagan, 1992), or, more technically, “phenomenal consciousness” and “access consciousness” (Block, 1995). It would be nice to be able to explain the exact nature of the connection between these two basic features of consciousness, assuming, of course, that the correlation between the feeling component and the knowledge component is more than mere accident. In particular, it would be nice to know whether one of these two constitutive aspects of consciousness is ontologically prior to the other, the latter being, in some sense, derivative of the former; or, if the two are mutually dependent and none is more basic than the other, to know how they contribute to each other’s structure and character. In short, it is desirable to be able to answer the question what have the knowledge aspect and the feeling aspect of consciousness to do with each other.

Alas, the problem is convoluted, admitting no simple answers. For the sake of making the present discussion manageable I shall narrow down the domain of inquiry by focusing on a question which can be framed in unidirectional terms, namely: What, if any, is the explanatory relevance of the knowledge aspect of consciousness (KAC) to the feeling aspect of consciousness (FAC)?

In other words, the question is how we should understand the role of sapience in the making, or shaping, of sentience.¹

This question can, in turn, be broken down into yet more specific components. In particular, it is useful to structure the discussion around two complementary issues. The first issue is whether or not it is possible to *derive* FAC from KAC. Clearly, if sentience is reducible to sapience then there is a very obvious sense in which KAC is explanatorily relevant with respect to FAC, namely, that the feeling aspect of consciousness is but a function of, or a specialization within, the knowledge aspect of consciousness. But if no such reduction is in the cards it becomes less clear what, if any, is the role of sapience in the making, and shaping, of sentience. Hence, in the eventuality that phenomenal consciousness cannot be derived from the purely informational components of consciousness, the next issue which confronts us is what alternative role might sapience still play in the explanation of sentience.

My goal is to articulate such a non-reductive alternative, namely, to explain in what sense sapience is still indispensable for a proper understanding of sentience even if we relinquish the hope (or the nightmare?) of reducing the latter to the former. Throughout the discussion, I shall assume that there is an important sense in which phenomenology is in the head, or, at any rate, the head and body. Those who are sympathetic to *phenomenal externalism* (see, e.g., Dretske 1995; Lycan 1996; Tye 1995), i.e., to the idea that the qualitative features of experience consist entirely in properties of the objects of experience (or, in other words, that there is nothing internal about the phenomenal character of conscious mental states) may find little to trouble them in what follows. To attempt to refute this influential position is something which I cannot do on the present occasion without sinning against bulk and thematic balance. There is enough sense, I believe, in exploring the relationships between sentience and sapience based on the traditional and still popular idea that consciousness presents us with an unflinchingly inner dimension of reality, leaving the question whether the belief in such an inner dimension is justified for another occasion.

FAC from KAC: Can Sentience be Reduced to Informational Access?

Those who are realists about phenomenal consciousness, and who take it to be a natural phenomenon, can agree at least on one thing: that consciousness

¹Here and elsewhere in this paper I use the term “sapience” as the informational correlate to the feeling aspect of consciousness (viz., to sentience, or phenomenal consciousness). Although this term is somewhat archaic its meaning reflects accurately the knowledge aspect of consciousness. As such, it has an advantage over more frequently used terms such as “awareness,” which are rife with phenomenal connotations. Thanks are due to Liam Dempsey for suggesting this term to me and for pointing to its earlier use by Feigl (1958).

as we know it, i.e., the kind of non-transitive consciousness with which we are all intimately familiar on a regular basis, is an ontological novelty. By describing it as an ontological novelty I mean simply this: that, as a natural kind, such consciousness did not exist from the very beginning of things. It came into being sometime during the course of cosmic evolution, indeed, by all accounts, relatively recently. If we were to go back in time two billion years, nothing here on Earth would be endowed with anything resembling our own states of consciousness, or those of other recently evolved intelligent species.

This much can be agreed upon not only by orthodox physicalists but also by panpsychists, neutral monists, and even absolute idealists. Thus, for example, panpsychists need not deny that consciousness as we know it is markedly distinct, qualitatively speaking, from the micro-phenomenal states which, they hypothesize, are enjoyed by unicellular organisms, molecules, etc. On the contrary, most panpsychists would agree that the differences between the micro-phenomenal and the macro-phenomenal levels are, in all likelihood, staggering.

However, attempts to describe the *nature* of this coming into being of macro-phenomenal consciousness quickly lead to wide disagreements. Physicalists typically accept a metaphysical picture according to which the antecedent physical conditions which gave rise to the evolution of macro-phenomenal consciousness, as well as the physical “building blocks” whose combination gives rise to tokens of experience, are ultimately devoid of subjectivity and sentience. In its default, aboriginal state, nature is utterly numb, lacking an interior and certainly lacking anything which remotely resembles phenomenal consciousness, or which could be considered a precursor of experience. To varying degrees, supporters of panpsychism, neutral monism, absolute idealism, and even certain versions of unorthodox physicalism all deny this basic assumption.

By affirming that nature’s default state is categorically objective and insentient, orthodox physicalists commit themselves inadvertently (to the extent that they are realists about consciousness, that is . . .) to the idea that phenomenal consciousness is a completely new ontological kind, categorically distinct from, and utterly discontinuous with, anything else in nature. Such a position is often referred to as *radical emergence* (see Seager and Allen–Hermanson, 2013; van Gulick, 2001), a view which Galen Strawson describes as holding that there is nothing “about the nature of the emerged-from (and nothing else) in virtue of which the emerger emerges as it does and is what it is” (2006, p. 15).²

²This commitment to radical emergence is inadvertent insofar as many (perhaps most) orthodox physicalists are firm believers in a physically reductionist explanation of consciousness. As reductionists, they would be very reluctant to align themselves with an idea whose flavor is reminiscent of the highly non-reductive doctrines of good old British emergentism. Yet, the blatant discontinuity between a “dead,” categorically insentient universe and the reality of subjective experience seems to leave the qualia realist with little choice but to affirm radical emergence.

It is possible for an orthodox physicalist to maintain that the emergence of sentience from the dim background of an utterly insentient world is a brute physical fact, the result of a dynamical configuration of processes (typically, neurophysiological ones) whose activation just happens to “switch the lights on” as it were. Indeed, I argue in the next section that such a scenario seems inevitable on the assumption that complete insentience is nature’s default state. However, as a matter of fact, the neural configurations which are being identified as the underpinning correlates of conscious experience are typically portrayed as psychologically meaningful. For example, in the works of such authors as Crick and Koch (1990), Dehaene and Naccache (2001), and Edelman (1992), the identified correlates of conscious experiences are all processes which are presumed to be responsible for large-scale integration and retrieval of information, often through attentional amplification. Likewise, in Damasio’s (1999) theory of core consciousness, experience is associated with meta-representations of the interactions between subject and world. In other words, the neural correlates of conscious experience are processes laden with cognitive significance corresponding to what was identified earlier as the knowledge aspect of consciousness.

That this is the case is not surprising given the robust correlation between FAC and KAC, or between sentience and sapience, but it leaves open the question just why the cognitive processes thereby identified are endowed with a phenomenal feel. In other words, there remains the question why must sapience be accompanied by sentience. Thus, while the explanatory burden may shift from the neurophysiological domain to the cognitive domain the gap in the explanation of phenomenal consciousness remains (see Chalmers, 1996). One way of approaching a solution to this problem is via what I call the *FAC-from-KAC* hypothesis, namely, the hypothesis that phenomenal consciousness is reducible to, or is a function of, representational relations between mental states which in themselves are not, or need not be, phenomenally conscious.

To recapitulate, the idea is that sentience is explainable in terms of sapience. In principle, it is possible to adopt a more modest stance: treating the correlation between sentience and sapience on a purely descriptive level. For instance, one could observe that certain forms of cognitive awareness (perhaps re-entrant signalling, or perhaps higher-order monitoring) are invariably accompanied by phenomenal consciousness and yet one can refrain from making the stronger claim that the latter results from, or is reducible to, the former. To argue in favour of the stronger thesis is to take an explanatory rather than a purely descriptive approach. This stronger thesis is, of course, more ambitious and therefore more exciting. Correspondingly, there is no shortage of bold theorists willing to pursue this explanatory project by articulating one variant or another of the *FAC-from-KAC* hypothesis. Below, I consider three different varieties of this ambitious agenda representing higher-order monitoring, self-representational, and thick specious present theories, respectively.

David Rosenthal, a leading advocate of the higher-order thought (HOT) account of consciousness, is a clear defender of the FAC-from-KAC hypothesis (see also Carruthers, 1996; Gennaro, 1996). On Rosenthal's (2002) view, a mental state *M* is conscious (i.e., non-transitively conscious) if and only if there is another mental state *M** such that *M** is an occurrent higher-order thought representing *M* as its object. Thus, it is in virtue of being represented by a HOT *M** (which in itself may be either conscious or unconscious, as the case may be) that *M* becomes conscious. Rosenthal is aware that it is, *prima facie*, far from obvious that the mere fact that *M* is represented by *M** should account for there being something it is like to be in state *M*. Nevertheless, he proceeds to defend just that, arguing that "being able to form intentional states about certain sensory qualities must somehow result in being able to experience those qualities consciously" (p. 413).

Rosenthal's justification of this bold claim is somewhat complicated but, in essence, it consists of the idea that our ability to be conscious of sensory qualities is contingent on our capacity for making appropriate conceptual discriminations, discriminations which are captured in the form of higher-order thoughts. The more able we are of making such conceptual discriminations, the greater the variety of sensory qualities we can experience. For example, "learning new concepts for our experiences of the gustatory and olfactory properties of wines typically leads to our being conscious of more fine-grained differences among the qualities of our sensory states" (ibid.). Conversely, he argues that the lesser the amount of classificatory HOTs at our disposal the duller and the more generic our experience becomes, such that peeling away all HOTs "would result, finally, in its no longer being like anything at all to have that sensation" (ibid.) The moral, then, is that HOTs are both necessary and sufficient for phenomenal consciousness.

Self-representational accounts of consciousness are often advanced in contrast to higher-order monitoring theories (see Kriegel, 2007), yet at least one defender of the self-representational view — Greg Janzen — shares Rosenthal's explicit endorsement of FAC-from-KAC. Janzen argues that "the phenomenal character of at least perceptual consciousness can be fully explained in terms of self-awareness, i.e., in terms of a low-level or 'implicit' self-awareness that is built into every conscious perceptual state" (2006, p. 44). On the self-representational view of consciousness, whose roots are traceable to Brentano (1874/1995), every conscious state is *bidirectional*: it is intentionally directed at (i.e., transitively conscious of) something other than itself, but in addition it also curls upon itself. Although largely *implicit* (owing to the fact that, normally, one's attention is focused on the intentional object rather than on the internal representational medium), such reflexive self-awareness is a constant presence, lurking in the background.

As mentioned above, Janzen argues that implicit self-awareness is literally constitutive of the phenomenal character of perceptual states and, possibly, of

many other conscious mental states. His account is deflationary, consisting, in essence, of the idea that the particular “what it’s like” of a perceptual state P consists of the particularities of the process of perceiving x, the perceptual object, while the self-awareness component guarantees that this “what it’s like?” of P is present to the cognitive subject. The result, allegedly, is that there is *something it is like for me*, the subject and the owner of P, to perceive x, and this, he argues, offers a coherent solution to the problem of phenomenal character.³

Finally, the FAC-from-KAC hypothesis is also espoused by Nicholas Humphrey (1992, 2000), as well as by Ralph Ellis and Natika Newton (Ellis and Newton, 2005) who, like Humphrey, pursue the subject from an action-oriented perspective in which the notion of a temporally thick present plays a prominent role. In both accounts the central idea seems to be that phenomenal consciousness, the raw feels of experience, results from the superposition of distinct temporal moments onto a unified thick specious present. This partial overlap between memorized past, occurrent present, and anticipated future allows for information to coalesce into coherent units of extended “thick” moments in which representations of past, present, and future are vividly accessed, and are recruited at the service of meaningful action guidance. According to these authors, phenomenal consciousness consists of nothing more than such happy coalescence of enactive representations.

Problems with FAC-from-KAC Reductionism

It is difficult not to feel, however, that something is amiss in all of these accounts. There is a lingering impression that they presuppose that which they seek to explain and that without such illicit presuppositions the explanations simply do not work. Consider first Rosenthal’s account. Sure, learning to make finer conceptual discriminations with respect to the taste of a wine, or to the sound of an oboe, allows for more refined experiences of the kind savoured by the connoisseur, but such finesse is the result of a process of training involving interaction between newly acquired knowledge and previous, more basic, experience. The connoisseur and the layperson experience taste, sound, or sight differently but they both operate within a space which is *already* richly experiential; the connoisseur is a specialist processor of experience, not the generator of experience out of insentience. Nevertheless, Rosenthal believes

³The idea that the explanation of phenomenal consciousness calls for a division of labor based on the conceptual distinction between a *qualitative* component (i.e., the something it is like aspect of seeing x) and a *subjective* component (the “for me” aspect of private experience) is due to Levine (2001). Kriegel (e.g., 2005, 2009) developed an influential self-representational account of consciousness in which this conceptual distinction plays a prominent role. However, unlike Janzen, Kriegel was never quite as adamant to declare that implicit self-awareness fully explains phenomenal character and recently he came to concede that a materialist reduction of phenomenal consciousness remains an elusive goal (Kriegel, 2011).

that the lesson from the connoisseur analogy is that there is a direct proportionality between the conceptually savvy and the experientially potent, a lesson which he then extrapolates to argue that peeling away all HOTs would eventually culminate in the absence of all phenomenal consciousness.

This suggestion, however, is highly non-intuitive. The idea that experience is constitutively dependent on conceptualization seems to put the cart before the horse: it would make the capacity to sense and feel which, by all accounts, is rather basic in evolutionary terms dependent on the abstract operations of conceptual thought — a much more sophisticated agent, and an evolutionary latecomer. Nor does the suggestion seem to fit with phenomenological data. Psychoactively induced experiences are, at times, remarkably rich, despite, and perhaps partly because of, the fact that the veil of rational classification and control is being lifted. Even more so, mystics of all ages consistently report that the cessation of all thought through meditation leads not to numbness and stupor but to the most intense and lofty experiences possible for humans (or think of the vividness, intensity, and freshness of the experiential reality of a child, which stands in sharp contrast to the child's lack of conceptual sophistication).

Viewed from a different angle, we may question not only the plausibility of Rosenthal's proposal but also its very intelligibility (cf. Goldman, 1993). For how could the mere fact that M is being represented by a higher-order thought M^* turn M into a phenomenally conscious state? Crucially, there is nothing in this scenario which implies an internal modification of M, let alone a radical modification of the sort which would be required in order to make it a locus of sentience. The relation of being represented by M^* is, insofar as M is concerned, an external relation, which implies, in turn, that whether or not the relation holds is something which has no effect, or need have no effect, on M's intrinsic qualities (compare: the fact that I represent the Eiffel tower as I think of it now induces no visible change in the tower itself, something more needs to be added if such a change is to be effected). How, then, could such a relation turn M from an insentient state (as per hypothesis) to a vehicle of sentience? The transition seems miraculous enough even if M^* did have a clear causal impact upon M, let alone when it has none!⁴

A careful reading of Rosenthal reveals that he also does not believe in such alchemy, arguing that being transitively conscious (*viz.*, aware) of a sensory state M does not change the properties of that state. Rather, the effect of the conceptual HOTs we apply to M is to “enable us to be conscious of sensory qualities we already had, but had not been conscious of” (2002, p. 413). But the problem refuses to go away: the idea that awareness of M's sensory qualities (courtesy

⁴It may be mentioned in passing that Rosenthal's conviction that phenomenal consciousness can be explained in strictly extrinsic terms is not shared by all HOT theorists. Gennaro (1996), for example, developed a HOT account which strives to accommodate the intrinsic character of non-transitive conscious mental states.

of the HOT M*) is the key for explaining sentience is illusory. For if we assume that both M and M* are wholly insentient then neither the qualities of M, nor the medium whereby they are represented in M*, are truly experiential. At best, all there is here is information in one wholly insentient physical state M* about the properties of another wholly insentient physical entity M (and note that the fact that the relevant properties are *sensory* properties makes no difference: for unless we assume a reading of “sensory properties” which illicitly introduces the reality of sentience then such properties would be just like any other properties of numb matter). In short, it is hard to see where in all of this phenomenal consciousness could ever be found. To suppose that informational liaisons between wholly insentient internal states could somehow result in there being sentience somewhere in the system is to expect the blind to successfully lead the blind, but, as the old saying goes, “when the blind lead the blind both shall fall into the ditch” (Matthew 15:14).

Nor do I think that the other theories mentioned above fare better. Janzen’s attempt to derive phenomenal character from implicit self-awareness ultimately faces the same problem. Self-representation is an internal relation, which is to say that the act of representation whereby P (a perceptual state) represents itself is part and parcel of P’s identity. This means that the “no difference” argument raised against Rosenthal cannot be raised here because self-representation does make a difference — P would not be quite the same if it was not self-representing. However, the problem is with the suggestion that the difference which this relation of self-representation brings about, or explains, is the difference between sentience and utter insentience. That is, if we assume a basic ontology of insentient matter, and if we assume that self-representation is the one crucial factor which delivers us onto the realm of sentience, we end up with absurdity.

The absurdity lies in the idea that the fact that a given state M curls upon itself, thereby instantiating an informational closed-loop, could somehow turn it into a locus of experience. For, how could M’s self-accessing be responsible for the transmutation? If the default assumption is that barring the self-referencing, M is just like any other physical state, which, per hypothesis, means an utterly insentient state, then there is “no one at home” to feel, sense, or be cognizant of the incoming information, and the fact that the information is self-originated, or self-effected, does nothing to change it. Or to put it differently, an insentient medium, or substance, cannot feel itself any more than it can feel any other thing — for it can feel nothing at all. Figuratively speaking, to suppose that self-representation can generate sentience from scratch is like supposing that a blind person can gain sight by staring at her own reflection in the mirror (see Levine, 2006, for an alternative argument against the idea that self-representation holds the key for solving the hard problem of consciousness).

The problem recurs in a different guise for theories that emphasize the role of the thick present moment. The coalescence of temporally differentiated rep-

representations onto a unified thick moment may constitute a significant step in solving the problem of explaining the possibility of a temporally extended awareness, that is, of an awareness which goes beyond the instantaneous moment of physical time, but it does nothing to usher in phenomenal consciousness. The contribution of the thick specious present is purely functional: it consists of enabling us to dwell on our experiences and to savour them, granting us access to an inner reality which would otherwise go unnoticed (see Ellis and Newton, 2005). But note that this could only work if we sneak in the implicit assumption that the reality which is thereby accessed is experiential through and through — clearly an illicit move for anyone who proclaims to explain the coming into being of phenomenal consciousness. For if we assume, as an orthodox physicalist should, that the relevant representations, call them M_P , M_N , and M_F (for past, present, and future), are decisively insentient, then it makes no sense at all to suppose that the operation of fastening them together such that there is a partial overlap between them could somehow result in there being a vividly phenomenal character to their overlap. Rather, we should expect to get just what we were constructing: a superposition of phenomenally vacuous states.

The moral of these consistent failures is that you cannot derive experience from information processing; you cannot get FAC from KAC. Representational access may serve to *transform* phenomenal character in myriad significant ways but it cannot generate sentience from scratch. If we start with the idea that the task is to derive experience from a decisively non-experiential realm, and that representations of one sort or another are our means to do so, we just end up producing more and more blind representations, blindly representing a blind world. Or in Levine's apt words: "it's just piling on more representations" (2006, p. 195).

If the FAC-from-KAC hypothesis is a dead-end street then the physicalist who is also a qualia realist and qualia internalist has to concede that the emergence of sentience out of an insentient world cannot be explained in terms of intra-representational accessibility. The alternative which seems to force itself upon her is that such emergence is a pure physical fact, devoid of an epistemic rationale. Somehow, certain complexly interacting organizations of matter, in particular neural activation patterns, manage to evoke experience as part of their activation, but we cannot explain such evocation in psychological terms.

This concession is somewhat disturbing given the ample evidence which suggests that sentience and sapience (FAC and KAC) go hand in hand and are intimately connected, but it may not be so worrisome if other sciences could step in and fill the explanatory gap. Yet, this hope, too, seems to be in vain. As Nagel (1974), Levine (1983), Chalmers (1995), and others have pointed out, the emergence of sentience against the background picture presupposed by orthodox physicalism

has the appearance of a hopelessly brute fact, and a high-level brute fact at that.⁵ It is no more explicable in physical, chemical, biological, or computational terms than it is in strictly psychological terms. The nagging question, “Why sentience?” remains just as vexing, no matter which scientific discipline we care to consult.

Indeed, the conceptual aporia to which the FAC-from-KAC hypothesis leads are but special exemplifications of the more general problem known as the explanatory gap (Levine, 1983), or the hard problem of consciousness (Chalmers, 1995). If this is the case, one may wonder why I have bothered to discuss at length the theories mentioned above only to end up returning to such a familiar point. The answer is that a firm understanding of the reasons behind the failure of the FAC-from-KAC hypothesis is necessary in order to motivate an alternative, non-reductive approach towards understanding the relations between sentience and sapience. In the remaining sections I present and defend the essentials of such an alternative.

Moderate Emergence and the Continuity Principle

As an alternative to the idea that sentience is derivable from epistemic liaisons between insentate mental states, I shall now pursue the idea that sentience and sapience go hand in hand — both co-evolve in correlation and none is more fundamental than the other. This parallelism is essentially in line with Chalmers’ (1996) *coherence principle*. However, Chalmers’ view has certain concomitant components which I am reluctant to accept, in particular: (a) his functionalist principle of organizational invariance according to which all functional isomorphs of a given conscious system S are experientially indiscernible from it; and (b) his analysis of awareness in terms of information processing, understood in the strictly syntactic sense of information theory. None of these latter elements is included in the present account.

Be that as it may, there is a more general issue which we must address before returning to the specifics of my proposal, to wit: What are the natural boundaries within which we should expect to find consciousness, complete with its correlative knowledge aspects and feeling aspects? Logically speaking, a parallelism between sentience and sapience implies only that they come together, but it tells us

⁵As Levine (1983) points out, certain physical facts simply *are* brute facts (at least from our human limited perspective). Yet, as he observes, the arbitrariness of sentience is particularly disturbing precisely because, unlike other brute facts such as the particular value of the gravitational constant, sentience is presumed to be a higher-level phenomenon, and higher-level phenomena are alleged to be *explicable* in terms of our theoretical understanding of the workings of lower-level phenomena. Chalmers’ plea for an ontologically *fundamental* theory of consciousness (1996, chap.8) can be seen as an attempt to correct this anomaly by situating psychophysical laws alongside the basic laws of physics.

nothing about their scope in the natural world. In particular, parallelism does not tell us whether consciousness is ontologically primordial or whether it is an evolutionary latecomer. However, given the earlier observation that the orthodox physicalist picture seems to render the reality of conscious experience irredeemably inexplicable, it stands to reason that if we wish to avoid this most unpleasant consequence we ought to open ourselves to the possibility that nature's default state may not be categorically exclusive of sentience and subjectivity. And since orthodox physicalism is implicitly committed to the idea that consciousness, if it exists at all, is a radical ontological emergent, a consistent alternative is likely to involve the idea that consciousness is a moderate emergent.

Moderate emergence is the thesis that there is continuity in cosmic evolution, such that if X emerges from background conditions $C_1 \dots C_n$ there must be something about $C_1 \dots C_n$ which, in principle, could render X 's emergence, and its unique characteristics, intelligible. In other words, the seeds of that which emerges must somehow be latent already within that from which it emerges (for historical precursors to this idea see Leibniz, 1704/1995a; and Peirce, 1892/1955; in particular the former's *law of continuity*, and the latter's concept of *synechism*). A special corollary of moderate emergence is that the emergence of creatures endowed with an internal dimension out of physical preconditions in which such a dimension is presumed completely absent is precluded on pain of violating the continuity principle. Thus, on the assumption that an intrinsic dimension is clearly manifest in the structure of our own consciousness, it follows that such a dimension must be an integral part of nature at all levels of organization.

Those who espouse this line of reasoning often make use of what Seager (2006) calls the *intrinsic nature argument*. Following the footsteps of Eddington (1928) and Russell (1927), they note that scientific explanations are limited to the structural–dispositional aspects of reality, leaving unaccounted the intrinsic nature of the entities which science purports to describe and explain. Thus, it is not so much that modern science denies the existence of such intrinsic natures, or qualities (let alone proves their inexistence) but, rather, that it ignores them. Combined with the claim that intrinsic qualities are a logical desideratum, and that consciousness provides us with an existential proof of their reality, it is then suggested that there is more to reality than what is currently subsumed under the conceptual umbrella of contemporary natural science, and that a more complete metaphysics will have to take into account the intrinsic nature of things (advocates of this line of reasoning include Chalmers, 1996; de Chardin, 1959; de Quincey, 2002; Lockwood, 1989; Maxwell, 1979; Nagel, 1979; Rosenberg, 2004; Seager, 2006; Shimony, 1997; Stoljar, 2001; Strawson, 2006).

The significance of the intrinsic nature argument in the current context lies in the fact that the argument provides elbow room for the scenario of moderate emergence. In a world where nature, in its primordial state, lacks intrinsic qualities, the emergence of sentience is destined to constitute a radical and

inexplicable ontological breach, but in a world where such qualities exist, and are the rule rather than the exception, the emergence of macro phenomenal consciousness may simply represent a natural outgrowth out of humbler origins of a similar kind.

However, there is much disagreement over the question how to interpret this kinship and the ontological continuity it implies. In particular, does having an intrinsic nature imply having sentience; or does it merely imply a certain potentiation towards sentience. Panpsychists take the continuity principle to imply that sentience scales all the way down, or, in other words, that experiencing subjects and their corresponding phenomenal properties are aboriginal. To be sure, the experiential reality of an atom, an organic molecule, a metazoan, or a primitive protozoan is very different from ours, but, the idea goes, they nevertheless enjoy certain experiences (present day defenders of panpsychism include, for example, de Quincey, 2002; Griffin, 1998; Rosenberg, 2004; Seager, 2006; Sprigge, 1983; Strawson, 2006). In contrast, others, whom we may identify as Russellian identity theorists, or panprotopsychists (see Chalmers, 2013), argue that the intrinsic natures of sufficiently primitive beings are wholly insentient and yet that such intrinsic natures are *proto*-phenomenal in the sense that, when properly combined, they instantiate experience in an intelligible manner (defenders of this view include, for example, Feigl, 1958; Lockwood, 1989; Maxwell, 1979; Pereboom 2011; Stoljar 2001). Finally, there are also those who endorse neutral monism in the tradition of Mach (1886/1959) and James (1912) and argue that the fundamental entities are phenomenal properties but that subjects capable of experiencing such properties emerge only at a later stage (for a recent defence of this view see Coleman, 2014).

The theoretical framework I shall present shortly is panpsychist. On this occasion, I make no systematic attempt to motivate panpsychism over and against the other positions just mentioned. Nor do I offer a defence of panpsychism against the charge that it faces a combination problem (the term is Seager's, 1995) which is every bit as hopeless as the hard problem of consciousness that haunts orthodox physicalism. Doubtless, these are issues which sympathizers of panpsychism must address and I have done, to some extent, elsewhere (Shani, 2010). However, on the present occasion my goal is not to validate panpsychism fair and square but, rather, to explore the modifications which such a view entails with regard to the relevance of awareness for the explanation of sentience. Ultimately, I argue that a panpsychist framework provides a more coherent picture of this explanatory relation than the one bequeathed upon us by physicalism. If my diagnosis is correct, then it ought to serve as yet another reason to resist orthodox physicalism while moving in the direction of assigning consciousness a greater role in the scheme of things. However, I must qualify myself by adding that even if my point is valid we cannot rule out, at this stage, the possibility that certain alternative monistic positions other than panpsychism — perhaps

neutral monism, or perhaps panprotopsychism — might be able to claim an equal degree of explanatory coherence with respect to the problem at hand.

More FAC through More KAC: Outlines of an Ampliative Approach

Having rejected the idea that epistemic liaisons are constitutive of phenomenal consciousness I propose instead that the contribution of sapience for the making of sentience is ampliative. On the ampliative model, informational liaisons modulate phenomenal character rather than generating it from scratch. Moreover, the modulation is enhancive, which is to say that there is a positive correlation between the representational complexity and power of a system and its phenomenological richness — an increase in one is conducive to an increase in the other. Correspondingly, from this perspective, the main explanatory challenge does not consist in explaining how phenomenal consciousness comes into being in the first place but, rather, in explaining how it *changes* as a function of changes in representational power.

The ampliative model can be characterized by five basic theses:

1. [Concomitance]: Every act of presentation, or of re-presentation, involves a subject in cognizance of an object, or a datum, which it presents, or represents, through a subjective medium, which reacts to the object, or datum, with feelings.⁶
2. [Endo-phenomenology]: Phenomenal character is an endogenous feature of the medium of representation; which is to say that even in the absence of stimulation the medium is still a locus of sentience.
3. [Transformation]: Presentational, or representational, acts operate on the medium as transformative agents, constraining and modulating the ever-present flow of experience.
4. [Correlation]: In general, there is a direct proportionality between the level of sophistication of a system's cognitive organization and the depth and variability of its phenomenal world.
5. [Enhancement]: Informational liaisons between representational states are often instrumental in enriching the structure and character of experience. In other words, the transformative effect of acts of awareness on the subjective medium is often in a qualitatively ascending direction.

[Concomitance] is reminiscent of Whitehead's (1929/1985) notion of *prehension*. The important point in the present context, however, is the concurrence between

⁶The distinction between *presentation* and *representation* parallels Searle's (1983), which means that it corresponds to the distinction between those situations in which the intentional object is present to one's senses and those in which the intentional object is not currently present and has to be re-presented in one's mind. For simplicity's sake, however, I will follow the common practice of using the term "representation" in a looser sense covering both presentations and re-presentations.

an endogeneously sentient substrate (viz., the medium) and the acts of awareness, or sensitivity, whereby the subject takes into account external data. In other words, the idea is that sentience and sapience are coextensive: a sentient system is simultaneously a system which exemplifies awareness to events within, and outside, itself; likewise, a system capable of genuine awareness is, concurrently, a sentient system.

[Endo-phenomenology] expresses the idea that the physical substrate which serves as a medium for occurrent conscious representations is inherently sentient. This means that in the absence of significant external stimulation such a medium maintains a relatively homogenous qualitative state (in a manner analogous to that of an energy field subject to no discernable local excitation), or, alternatively, that it generates its own activation patterns, perhaps subject to chance events. External stimuli create stirs, or waves, on the surface of this “ocean” of spontaneous activity which in turn effect further transformations down the line, inducing changes in the patterns of organization that characterize the medium at the time. Thus, the structure of the medium (of which phenomenal tone is an essential aspect) is responsive to the structure of the environments with which the subject interacts.

[Transformation] serves to emphasize that the ever changing flow of experience is modulated by representations. That is, both representations of the outside world (by way of anticipation, perception, memory, or imagination) and of the self (i.e., representations of activities within the system or of the manner in which the system is influenced by external encounters) induce changes in the structure and course of the system’s internal experiential flow, leading to consequent representations and consequent process modulations down the line.

[Correlation] stresses a direct proportionality between the representational complexity exemplified by a cognitive agent and the phenomenal riches which the agent enjoys (or can enjoy). To use an extreme example, there is little reason to doubt that the mental reality of an orangutan is considerably richer than that of a jelly fish, not only on account of cognitive sophistication but also in terms of phenomenal variability and depth. These differences are indicative of a general rule, applicable throughout the animate world: the higher we go up the evolutionary ladder we find greater riches both in terms of sapience and in terms of sentience (at the same time, we must guard against the tendency to downplay the emotional and cognitive sophistication of relatively simple creatures, or to ignore their uniqueness). Conversely, if experience is something which even inanimate entities are presumed to possess, it is natural to expect this general rule to continue to hold all the way down so that, in the words of Teilhard de Chardin, “[r]efracted rearwards along the course of [cosmic] evolution, consciousness displays itself qualitatively as a spectrum of shifting hints whose lower terms are lost in the night” (1959, p. 59, italics in the original).

Finally, [Enhancement] drives home the point that the more a system is capable of accessing and processing its own experiences the more it is capable of having experiences of novel kinds, thereby intensifying its experiential reality. Consequently, [enhancement] constitutes the one aspect of the ampliative model most relevant for the present discussion, and the one on which I focus henceforth.

That [Correlation] obtains is something which calls for an explanation, and the best explanation seems to be that our two complementary dimensions of conscious experience — sentience and sapience — are mutually reinforcing. The gist of the idea is that greater representational power is conducive to greater variability and intensity in a creature's phenomenal life; and collaterally, an increase in the scope and intensity of phenomenal expression augments the capacity for representational classification, leading to novel and more articulated forms of awareness, and of action-guidance through awareness. Mutual reinforcement is, of course, a bidirectional relation but in line with my earlier resolve I focus here on the amplificatory effect of incremental awareness on phenomenal consciousness.

For illustrative purposes, imagine a creature which we may call Primo. Primo is a blobby little creature whose protoplasmic interior manifests a minimal degree of internal organization. It is, however, sentient. It detects certain chemical gradients, and reacts to light, heat, and mechanical contact. These environmental interactions translate to internal events one aspect of which is that they create ripples in Primo's drearily shallow endo-phenomenological pond. Some of these ripples are recurrent and systematic enough to play a role in guiding Primo's behaviour. For the most part, however, ripples (whether spontaneous or externally induced) appear across the pond only to disappear quickly without leaving visibly recognizable traces. Yet, Primo is a special creature. It goes through a developmental catastrophe after which it changes quickly and radically. It grows in size; its internal milieu differentiates to various compartments, giving rise to a multitude of well-coordinated organelles, cells, tissues, and organs; it also grows external organs, some specialized for locomotion and object manipulation, some for the detection of information; it even grows an impressively dense ganglia full of interconnected nerve cells which enable it to integrate information from its newly grown perceptual and motor organs (as well as bodily surface) with information from its newly grown internal milieu, to make records of such informational confluence, to recall traces of those records, and to use all of this in guiding the activities of its monstrously changed self. In short, Primo is a one-in-all evolutionary freak.

Clearly, we should expect post-catastrophic Primo, call it Primo₂, to enjoy a richer phenomenal reality than its pre-catastrophic self Primo₁, but the question is why. One explanation, which is in line with much of contemporary thinking

about panpsychism, is that this has to do with the fact that Primo_1 is but a tiny micro-organism whose phenomenal field is limited to micro-experiences with micro-phenomenal properties, whereas Primo_2 is a multi-cellular organism whose phenomenal field combines the phenomenal fields of its micro-components, giving rise to a rich tapestry of macro-experiences endowed with macro-phenomenal properties. Now, whether or not such a combination story makes sense (recall the combination problem) there is no denial that bulk is a factor in the differences between Primo_1 and Primo_2 . To use aquatic metaphors, if Primo_1 's endo-phenomenological space is a shallow pond then Primo_2 's is a vast ocean, and, as we know, it takes an ocean to manifest certain wave patterns.

But, of course, this is only part of the story. Patterns of ripples and waves (our analogy for experiences) depend on other factors: wind currents, the moon, volcanic activity, local movements of vessels, objects, and animals, the throwing of stones, even artificial wave generators. In the end, what matters are the *patterns* of disturbance generated and bulk is, at best, only a necessary condition for that. To go back to the thought experiment, the moral to take home is that if we wish to explain the spectacular differences between the phenomenal realities of Primo_1 and Primo_2 we must look for the formative agency, the "wave generator" responsible for creating such vast differences in the patterns of disturbance characteristic of the respective endo-phenomenological media of these creatures.

To continue this idea, I think that the fact that the experiential life of Primo_2 is so much richer than that of Primo_1 depends crucially on the enormous differences in their degrees of internal *organization*. Primo_2 is a complexly organized creature capable of constraining, directing, and regulating the flow of energy, and the distribution of work, throughout itself in multitudinous ways unavailable to Primo_1 . This increased capacity for self-governance is, I suggest, the formative agency we need to look at.

Clearly, the development of more powerful representational capacities is an aspect of advanced self-governance. It enables improved process coordination, anticipation, action-selection, and much more (for further discussion of the connection between representation and self-governance see Bickhard, 2000; Clark, 1995; Collier and Hooker, 1999; Kauffman, 2000; Pezzulo, 2011; Shani, 2006). The ability to know more, with better resolution, in greater detail, and with greater depth and scope allows for the possibility of more refined self-governance and opens up new horizons for practicing novel forms of interaction and self-conduct. This much is evident, but the reason I mention it here is the *formative* influence on the qualities of experience. Unlike Primo_1 , Primo_2 enjoys vast representational resources: a wide spectrum of sensory, somato-sensory, motor, and visceral differentiations, which enable the formation of a plurality of perceptual and other presentational states; the ability to memorize, and to re-enact memorized representations; a capacity to form prospective representations anticipating

future conditions; powerful means for processing and integrating cognitive information; the ability to monitor its own internal events and to interact with them, using higher-order states, etc.

My point is that these functional aspects serve to augment and enrich experience. Primo_1 has a very limited access to the world as well as to its own internal conditions. In contrast, Primo_2 has a broader, deeper, and more articulated access to the world around it, and in addition it also has far more sophisticated ways to access its inner reality. These windows on the world and on the self are really operations which induce novel and ever more refined disturbance patterns in the creature's endo-phenomenological space, thereby enriching the landscape, or texture, of that space.

It is easy to see how a greater ability for making perceptual discriminations augments one's phenomenal world — it creates more experiences, and more shades of experience. But so is the case with the capacity to operate upon one's own representations, to experience one's own experiences, as it were. When an image is recalled and re-lived in the light of present experience, when one's own feelings are addressed, when a connection between different elements in one's experience is discerned and illumined by awareness; in short, whenever consciousness loops upon itself and the flow of experience becomes an object of experience, new *types* of experience emerge which were not available before. This recursive process is seemingly boundless — there are always novel and more refined experiences to be distilled, provided that the distillery (*viz.*, the system's organization) is up for the task.

Thus, the difference in Primo's experiential life before and after the morphogenetic mutation is, in large part, a difference in the capacity of its bodily organization to whip the waters of consciousness into shape. This, then, is the idea behind the ampliative model: that an increase in the capacity for information processing and access transforms and enriches the texture of one's phenomenal life without, however, being responsible for the fact that there is experience in the first place.

Putting the Ampliative Approach to Work

Above, I criticized theories committed to the FAC-from-KAC hypothesis for being caught in an explanatory cul-de-sac. I now proceed to show that once we translate these theories from their natural reductive setting to the non-reductive landscape delineated by the ampliative approach, we can restore coherence to some of their more attractive features — although, naturally, this process involves a reinterpretation of the meaning and scope of these theories.

Recall, first, Rosenthal's HOT-based account of phenomenal consciousness. I argued that the idea that a mental state M could become phenomenally conscious in virtue of being represented by a higher-order thought M^* defies sense,

but now look at the situation from the perspective of the ampliative model. The model assumes that the higher-order monitoring process $M^* \rightarrow M$ explains neither why M is phenomenally conscious, nor why M^* is. However, it predicts that such higher-order monitoring will result, typically, in experiences which are novel in kind, i.e., experiences of a kind whose existence is contingent on this very process. But where should we look for such experiences? Clearly, the object-level, the level at which M itself is located, would be the wrong place to look for such emergent phenomenology since, as mentioned earlier, the mere fact that M^* represents M does not imply any modifications in M itself (unless, of course, the higher-order monitoring process is an intervening one). Rather, it is to the meta-representational level, M^* 's level, that we should turn.

At the meta-level (or levels), we find mental states whose intentional objects are other mental states, and which represent qualities of those object-level states even as the latter represent qualities of the environment, or of the body (Bickhard, 2005). Thus, the features represented at the meta-level are different than the ones represented at the object-level. In particular, they may include such abstract elements as relations among the contents of object-level mental states, relations such as causality, similarity, ordering, matching, etc. (see, for example Barsalou, 1999; Chapman and Agre, 1986; Pezzulo, 2011). And awareness of such relations (perhaps courtesy of levels of representation higher-up the hierarchy) engages novel experiences, including an experiential type which is crucial for the discussion below, namely, the experience of *feeling oneself as an integrated experiential subject*. Thus, there is a grain of truth in Rosenthal's claim that higher-order cognitive processes are conducive to more refined phenomenologies; it is just that we can't expect meta-cognition to be the ultimate explanation of the reality of phenomenal consciousness.

Nor can we expect same-level self-awareness to carry the task. Above, I argued that self-awareness is of little help as long as we continue to assume that that which is being accessed, or in this case that which accesses itself, is inherently insentient. For if a mental state M is realized in an utterly insentient medium then it can neither be a locus of experience, nor can reflexive access grant it acquaintance with its own (non-existent) "experiential content." However, as soon as we change our default axiom to one in which experience occupies a fundamental place in nature, things begin to make better sense. First, hypothesis M is now realized in a medium which is inherently sentient, hence we should have no problem understanding how it could be a locus of experiential content. Second, we can now begin to make sense of the import of self-awareness: for if M , an inherently sentient state, represents itself, we should expect such access to yield conscious awareness of the experiential content enfolded in M . Such awareness would take the form of experiential acquaintance with M 's base-level experiential content, where both the base-level experiential content and the higher-level acquaintance with that content are

complementary aspects of M's phenomenal portrait. Supporters of the self-representational view are correct to emphasize the significance of self-awareness, for, clearly, the ability to be aware of one's inner reality is a crucial ingredient of consciousness *as we know it*, that is, as we find it in the structure of human phenomenology (see, for example, Zahavi, 2005 p. 24). Their mistake lies in the failure to realize that the contribution of self-awareness is intelligible only against a background which is already sentient.⁷

Lastly, consider the specious present account. The ampliative model is rather congenial to the idea that the ability to experience past, present, and future in a single "specious moment" is a significant landmark of consciousness. However, a careful scrutiny of this idea reveals that its real value lies not in the fact that it explains the transmutation of utterly insentient representations into a single complex locus of sentience, for this it does not do. Rather, the real contribution of the specious present with respect to phenomenal consciousness lies in the fact that it equips cognitive agents with a temporal window wide enough to enable us to become acquainted (i.e., experientially acquainted!) with our ongoing experiential flow.⁸ As such, it constitutes a major step in the discovery that we are enduring subjects of experience, but it does not explain the emergence of experience from the non-experiential.

Time and again, then, we see that the real contribution of sapience to the explanation of phenomenal consciousness is transformative and ampliative: it is instrumental in explaining how novel qualitative types of experience emerge atop other, more basic ones. Yet, no matter how hard we search, never do we find a single instance in which epistemic liaisons generate sentience from scratch. More from less everywhere, but nowhere is there something from nothing.

The Discovery of Experience: A Layered View of the Evolution of Consciousness

While epistemic access does not, and cannot, beget phenomenal consciousness, it plays a crucial role in explaining an important stage in the evolution of conscious experience, namely, that stage wherein a system acquires the capacity to *experience itself as an integrated subject of experience*. In other words, there is, indeed, a sense in which sapience is indispensable for an explanation of sentience but

⁷It might be the case that some supporters of the self-representational view (especially within the phenomenological tradition) are not committed to the constitutive approach and may even be sympathetic to the point I am making, yet I'm unfamiliar with any clear admission of this point. Thus, whether the point is denied, or whether the issue is insufficiently clarified, the overall impression is that advocates of the self-representational view succumb to the FAC-from-KAC fallacy.

⁸This idea is stated rather clearly by Ellis and Newton (2005) except that they fail to see with sufficient clarity that, as a matter of fact, what they explain is not our capacity to experience in the first place but, rather, our capacity to experience our own experiences!

this sense is limited to a higher-order form of phenomenal consciousness, consisting of sustained experiential acquaintance with one's own experiential life. Thus, it is not experience as such which is contingent on robust epistemic liaisons between inner mental states, but, rather, something different and more intricate, namely, the subjective discovery of the fact that one is the owner of an inner experiential realm! Such a sense of ownership over one's experiential domain ought not to be confused with full-blown *self*-consciousness since, in its most basic form, it implies neither possession of the concept of self, nor of a temporally extended ("autobiographical") sense of self (see below). Nevertheless, the capacity to sense the flow of one's experience as an integrated subjective arena is a precursor of mature self-consciousness, as well as of other high-level manifestations of reflective consciousness.

Prima facie, the idea that acquaintance with one's own experiential reality is an emergent phenomenon is provocative and even paradoxical since it can be easily interpreted as suggesting that below that level of emergence are creatures (or entities) which, although phenomenally conscious, are completely unaware of their inner realities. Now, this is a strange proposition. It is widely held that the very condition of being in a phenomenally conscious state implies awareness of that state (see, e.g., Chalmers, 1996; Kriegel, 2005). If so, then there can be no such thing as a phenomenally conscious yet wholly unnoticed, or unannounced, mental state, and this, in turn, cuts against the idea that it is possible for a creature to be phenomenally conscious without being in the least aware of its inner experiential flow.

In response, I should stress that I do not claim that it is possible to be phenomenally conscious without exemplifying *any* degree of awareness whatsoever. Nor does such a result follow from my analysis. Rather, my claim is more qualified, namely, that a certain degree of (emergent) epistemic access is a prerequisite for a certain degree of reflexive acquaintance with one's experiential flow, and that to the extent that such a degree of epistemic access is compromised it also compromises one's familiarity with one's underlying phenomenal reality. Or to put it in more concrete terms, the point I am making about higher-order experiential access to one's own experiences is that such access is a prerequisite for a *stable integrated acquaintance with one's inner reality, experienced as one's own* — I make no claim to the effect that a system which lacks such higher-order access to its own experiences lacks any kind of sensitivity whatsoever to its inner reality.

To illustrate the idea think first of a creature who is as simple as Primo_1 , or perhaps even simpler. As mentioned before, such a creature would be subject to various kinds of experiences, various kinds of disturbances to the endogenous oscillatory patterns of its endo-phenomenology. Some disturbances would be powerful, systemic, or significant enough to consume the creature's attention (however diffusive or automated it may be), or to trigger adaptive responses.

However, per hypothesis, such a creature is extremely primitive: its capacity to retain traces of its past experiences and to re-enact such memories, as well as its capacity to form anticipatory images of future events, are so rudimentary that it is virtually confined to an eternal phenomenal present. Moreover, and crucially for our present concern, the creature lacks the ability to form higher-order representations which would enable it to reflect back on its subjective experiential flow, to integrate its base-level experiences into a meaningful (and accessible) whole, and to appreciate the qualities and interrelations of such experiences. Thus, it has no means of discerning lasting relationships between classes of internal events, and of consciously recognizing their significance. In short, the poor creature's phenomenal world is both punctated and flat: fleeting experiences come and go like actors on stage but the agent having those experiences is just too amorphous to maintain a clear sense of ownership over the show.

Admittedly, this scenario is somewhat extreme, but we need not be afraid to think in extreme terms when probing into the possibility of sentience below the level of multi-cellular organisms, let alone below the bar of biological existence. The point I wish to stress is that a hypothetical creature of the sort just imagined illustrates the possibility of having a phenomenal life while, at the same time, being almost totally unaware of the fact that one has such a life. Knowledge of the fact that one is the owner of a private domain, acquaintance with the secret of one's own subjectivity, requires much more. It is like a mystery into which only some are initiated (and then, only partially and gradually) namely, those creatures whose organizational features enable them to loop over themselves, turning their experiential flow into an object of experience and observing the privacy of their inner world unfolds.

A prominent contemporary advocate of the idea that a sense of ownership over one's subjective reality is an emergent construction contingent on higher-order modes of access is Antonio Damasio (1999). Damasio's view of consciousness and selfhood is a layered view in which cognition builds upon emotion, which in turn builds upon homeostatic regulation (see also Dempsey and Shani, 2013; Watt, 2004). At the basis of his analysis is a layer he calls the *proto-self*, which is "a coherent collection of neural patterns which map, moment by moment, the state of the physical structure of the organism in its many dimensions" (p. 154). As the organism interacts with items in its environment it forms images of these items, and the *proto-self* undergoes modifications in response to these images, which, in turn, give rise to emotional reactions, and to mental images, or "feelings," responsive to such reactions. Now, according to Damasio, these patterns of interrelationships between images of the items with which the system interacts and the corresponding modifications to the system's *proto-self* are captured by higher-order representations recording the manner of change. In turn, these higher-order representations are responsible for the construction of a higher-layer of selfhood which Damasio calls *core consciousness*, and that

consists of a momentary sense of self in the act of being internally modified. Core consciousness, Damasio argues, provides for a sense of ownership over one's own inner reality, albeit a rather basic sense which is neither temporally extended (autobiographical), nor one which depends on verbal ability or on moral sense. It is a sense of ownership over one's private world which, presumably, many animals are capable of exemplifying but which requires a level of sophistication far beyond that of a creature like *Primo*₁.

What makes Damasio's account apt for the present discussion is the fact that it provides an empirical model (which incidentally also hints at the significance) of the ontological partition between (i) being an abode of subjective experience; and (ii) being capable of experiencing one's own subjective domain as *one's own*. Unfortunately, Damasio is somewhat unclear as to whether "the feeling of what happens" effected by his higher-order representations is a feeling of lower-order experiences or, rather, of utterly insentient occurrences. While only the first interpretation corresponds to the view I advance here, I believe that his work helps clarify the claim that an ontological partition of the sort just mentioned is a potentially important one.⁹

Such a partition was emphasized by Leibniz in his distinction between perception and apperception — the former being an inner state of the monad representing external things, while the latter consists of reflective knowledge of this inner state and is the prerogative of true minds (see Leibniz, 1714/1995b, 1714/1995c). It is also echoed in Whitehead's distinction between prehension and consciousness and in his claim that "consciousness presupposes experience" (1929/1985, p. 53). Leibniz complained that this failure to appreciate that not all perceptions are apperceived, or to put it in our terms, not all states of experience are objects of experience, led the Cartesians to their notorious belief "that [rational] minds alone are monads, and that there are no souls in animals, and still less other *principles of life*" (1714/1995b, p.197).

Few of us today would adhere to such stark Cartesianism yet, clearly, our collective legacy is much more Cartesian than Leibnizian. Doubtlessly, this legacy has some role to play in the reluctance of many to give any credibility to the possibility of hidden grades of consciousness scaling down throughout the whole of nature. Moreover, the fact that reflective awareness plays such a prominent

⁹On the one hand, Damasio describes his higher-order representations as "feelings of feelings" in a way which suggests that his higher-order representations are higher-order experiences representing lower-order experiences. On the other hand, since core consciousness materializes only at the level of higher-order representation his account leaves lower-order "feelings" below the bar of consciousness. It is interesting to note, however, that Damasio's theory of core consciousness is not aimed at explaining phenomenal consciousness as such but, rather, the sense of familiarity, identification, and ownership, which one feels with respect to one's inner reality (hence, the *scire* or knowledge connotation of "consciousness"). Thus, I believe that in essence his theory is consistent with an interpretation according to which the "feeling of what happens" is the sensing of one's own experiential flow.

role in the making of human consciousness makes it doubly difficult for us to think or to imagine clearly the possibility of conscious forms lacking all but the most primitive forms of awareness, or, to put it in Leibniz' terms, to conceive of monads which are not minds, or rational souls. Leibniz's warning against the potential detrimental consequences of failing to see behind the veil of our own phenomenology remains as relevant today as it was in his day.

References

- Barsalou, L. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577–600.
- Bickhard, M.H. (2000). Autonomy, function and representation. *Communication and cognition — artificial intelligence* [Special issue on: The contribution of artificial life and the sciences of complexity to the understanding of autonomous systems], 17(3–4), 111–131.
- Bickhard, M.H. (2005). Consciousness and reflective consciousness. *Philosophical Psychology*, 18, 205–218.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18, 227–247.
- Brentano, F. (1995). *Psychology from an empirical standpoint*. London: Routledge. (originally published 1874)
- Carruthers, P. (1996). *Language, thought and consciousness*. Cambridge: Cambridge University Press.
- Chalmers, D.J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2, 200–219.
- Chalmers, D. (1996). *The conscious mind: In search of a fundamental theory*. New York: Oxford University Press.
- Chalmers, D. (2013). *Panpsychism and panprotopsychism*. The Amherst Lecture in Philosophy 8: 1–35. <http://www.amherstlecture.org/chalmers2013/>
- Chapman, D., and Agre, P. (1986). Abstract reasoning as emergent from concrete activity. In M. P. Georgeff and A.L. Lansky (Eds.), *Reasoning about actions and plans — Proceedings of the 1986 workshop* (pp. 411–424). San Mateo, California: Morgan Kaufmann.
- Clark, A. (1995). Moving minds: Situating content in the service of real-time success. *Philosophical Perspectives*, 9, 89–104.
- Coleman, S. (2014). The real combination problem: Panpsychism, micro-subjects, and emergence. *Erkenntnis*, 79, 19–44.
- Collier, J.D., and Hooker, C.A. (1999). Complexly organized dynamical systems. *Open Systems and Information Dynamics*, 6, 241–302.
- Crick, F.H.C., and Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, 2, 263–275.
- Damasio, A. (1999). *The feeling of what happens*. New York: Harcourt, Brace, and Co.
- Dehaene, S., and Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, 79, 1–37.
- Dempsey L.P., and Shani, I. (2013). Stressing the flesh: In defence of strong embodied cognition. *Philosophy and Phenomenological Research*, LXXXVI, 590–617.
- de Quincey, C. (2002). *Radical nature: Discovering the soul of matter*. Montpelier, Vermont: Invisible Cities Press.
- Dretske, F. (1995). *Naturalizing the mind*. Cambridge, Massachusetts: MIT Press.
- Eddington, A. (1928). *The nature of the physical world*. New York: MacMillan.
- Edelman, G. (1992). *Bright air, brilliant fire*. New York: Basic Books.
- Ellis, R.D., and Newton, N. (2005). The unity of consciousness: An enactivist approach. *Journal of Mind and Behavior*, 26, 225–280.
- Feigl, H. (1958). The “mental” and the “physical.” In H. Feigl, M. Scriven, and G. Maxwell (Eds.), *Concepts, theories, and the mind body problem*. Minnesota Studies in the Philosophy of Science, Volume 2. Minneapolis: University of Minnesota Press.
- Flanagan, O. (1992). *Consciousness reconsidered*. Cambridge, Massachusetts: MIT Press.

- Gennaro, R. (1996). *Consciousness and self-consciousness: A defense of the higher-order thought theory of consciousness*. Amsterdam: John Benjamins.
- Goldman, A. (1993). Consciousness, folk psychology, and cognitive science. *Consciousness and Cognition*, 2, 364–382.
- Griffin, D.R. (1998). *Unsnarling the world-knot: Consciousness, freedom, and the mind–body problem*. Eugene, Oregon: Wipf and Stock.
- Humphrey, N. (1992). *A history of mind: Evolution and the birth of consciousness*. New York: Simon and Schuster.
- Humphrey, N. (2000). How to solve the mind–body problem. *Journal of Consciousness Studies*, 7, 5–20.
- James, W. (1912). *Essays in radical empiricism*. New York: Dover.
- Janzen, G. (2006). Phenomenal character as implicit self-awareness. *Journal of Consciousness Studies*, 30, 44–73.
- Kauffman, S.A. (2000). *Investigations*. New York: Oxford University Press.
- Kriegel, U. (2005). Naturalizing subjective character. *Philosophy and Phenomenological Research*, 71, 23–57.
- Kriegel, U. (2007). Philosophical theories of consciousness: Contemporary western perspectives. In M. Moscovitch, E. Thompson, and P.D. Zelazo (Eds.), *Cambridge handbook of consciousness* (pp. 35–66). Cambridge: Cambridge University Press.
- Kriegel, U. (2009). *Subjective consciousness: A self-representational theory*. Oxford: Oxford University Press.
- Kriegel, U. (2011). Self-representationalism and the explanatory gap. In J. Liu and J. Perry (Eds.), *Consciousness and the self: New essays* (pp. 51–75). Cambridge: Cambridge University Press.
- Leibniz, G.W. (1995a). New essays on the human understanding. In G.H.R. Parkinson (Ed.), *Philosophical writings* (pp. 148–171). (original work published 1704)
- Leibniz, G.W. (1995b). Monadology. In G.H.R. Parkinson (Ed.), *Philosophical writings* (pp. 179–194). (original work published 1714)
- Leibniz, G.W. (1995c). Principles of nature and grace. In G.H.R. Parkinson (Ed.), *Philosophical writings* (pp. 195–204). (original work published 1714)
- Levine, J. (1983). Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly*, 64, 354–361.
- Levine, J. (2001). *Purple haze: The puzzle of consciousness*. New York: Oxford University Press.
- Levine, J. (2006). Conscious awareness and (self) representation. In U. Kriegel and K. Williford (Eds.), *Self-representational approaches to consciousness* (pp. 173–198). Cambridge, Massachusetts: MIT Press.
- Lockwood, M. (1989). *Mind, brain, and the quantum: The compound 'I'*. Cambridge, Massachusetts: Basil Blackwell.
- Luu, P., and Tucker, D.M. (2004). Self-regulation by the medial frontal cortex: Limbic representation of motive set points. In M., Beauregard (Ed.), *Consciousness, emotional self-regulation, and the brain* (pp. 123–161). Amsterdam: John Benjamins.
- Lycan, W. (1996). *Consciousness and experience*. Cambridge, Massachusetts: MIT Press.
- Mach, E. (1959). *The analysis of sensations and the relation of the physical to the psychical [Die Analyse der Empfindungen und das Verhältnis des Physischen zum Psychischen; C.M. Williams, Trans.]*. New York: Dover. (Originally published 1886)
- Maxwell, G. (1979). Rigid designators and mind–brain identity. In C.W. Savage (Ed.), *Minnesota studies in philosophy of science* (Volume 9, pp. 365–403). Minneapolis: University of Minnesota Press.
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83, 435–450.
- Nagel, T. (1979). Panpsychism. In T. Nagel, *Mortal questions* (pp. 181–195). Cambridge: Cambridge University Press.
- Panksepp, J. (2005). Affective consciousness: Core emotional feelings in animals and humans. *Consciousness and Cognition*, 14, 30–80.
- Peirce, C.S. (1955). The law of mind. In J. Buchler (Ed.), *Philosophical writings of Peirce* (pp. 339–353). New York: Dover. (originally published 1892)
- Pereboom, D. (2011). *Consciousness and the prospects of physicalism*. New York: Oxford University Press.

- Pezzulo, G. (2011). Grounding procedural and declarative knowledge in sensorimotor anticipation. *Mind and Language*, 26, 78–114.
- Rosenberg, G. (2004). *A place for consciousness: Probing the deep structure of the natural world*. New York: Oxford University Press.
- Rosenthal, D.M. (2002). Explaining consciousness. In D.J. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings* (pp. 406–421). New York: Oxford University Press.
- Russell, B. (1927). *The analysis of matter*. London: Kegan Paul.
- Seager, W. (1995). Consciousness, information, and panpsychism. *Journal of Consciousness Studies*, 2, 272–288.
- Seager, W. (2006). The intrinsic nature argument for panpsychism. *Journal of Consciousness Studies*, 13, 129–145.
- Seager W., and Allen–Hermanson, S. (2013). Panpsychism. In E.N. Zalta (Ed.), *Stanford encyclopaedia of philosophy*. <http://plato.stanford.edu/archives/fall2013/entries/panpsychism/>
- Searle, J.R. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge, Massachusetts: MIT Press.
- Shani, I. (2006). Narcissistic sensations and intentional directedness: How second-order cybernetics helps dissolve the tension between the egocentric character of sensory information and the (seemingly) world-centered character of cognitive representation. *Cybernetics and Human Knowing*, 13, 87–110.
- Shani, I. (2010). Mind stuffed with red herrings: Why William James' critique of the mind-stuff theory does not substantiate a combination problem for panpsychism. *Acta Analytica*, 25, 413–434.
- Shimony, A. (1997). On mentality, quantum mechanics and the actualization of potentialities. In R. Penrose (with A. Shimony, N. Cartwright, and S. Hawking), *The large, the small, and the human mind* (pp. 144–159). New York: Cambridge University Press.
- Skrbina, D. (2006). Beyond Descartes: Panpsychism revisited. *Axiomathes*, 16, 387–423.
- Sprigge, T.L.S. (1983). *The vindication of absolute idealism*. Edinburgh: Edinburgh University Press.
- Stoljar, D. (2001). Two conceptions of the physical. *Philosophy and Phenomenological Research*, 62, 253–281.
- Strawson, G. (2006). Realistic monism: Why physicalism entails panpsychism. *Journal of Consciousness Studies*, 13, 3–31.
- Teilhard de Chardin, P. (1959). *The phenomenon of men*. London: Collins.
- Tye, M. (1995). *Ten problems of consciousness: A representational theory of the phenomenal mind*. Cambridge, Massachusetts: MIT Press.
- van Gulick, R. (2001). Reduction, emergence, and other recent options on the mind–body problem. *Journal of Consciousness Studies*, 8, 1–34.
- van Gulick, R. (2014). Consciousness. In E.N. Zalta (Ed.), *Stanford encyclopaedia of philosophy*. <http://plato.stanford.edu/archives/spr2014/entries/consciousness/>
- Watt, D.F. (2004). Consciousness, emotional self-regulation and the brain: Review article. *Journal of Consciousness Studies*, 11, 77–82.
- Whitehead, A.N. (1985). *Process and reality* [corrected edition]. D.R. Griffin and D.W. Sherburne (Eds.). New York: The Free Press. (originally published 1929)
- Zahavi, D. (2005). *Subjectivity and selfhood: Investigating the first-person perspective*. Cambridge, Massachusetts: MIT Press.