

Some Philosophical and Empirical Implications of the Fringe

BRUCE MANGAN

Institute of Cognitive Studies, University of California, Berkeley, California 94720

I will respond to the commentaries in the order I received them.

Alison Gopnik's "psychopsychology" is an interesting coinage. Probably many new terms will be needed as the scientific study of consciousness matures. Elsewhere (Mangan, 1991), I use "convergent phenomenology" to refer to a method that employs converging evidence from *as many sources as possible* to raise or lower the probability of a phenomenological claim. (This is quite different from the far narrower "Pure" phenomenology deriving from Husserl or Dennett's anti-realist "heterophenomenology.") I suspect that during the next decade many disciplines (e.g., linguistics, anthropology, sociology, aesthetics, religious studies, neurophysiology, jurisprudence) will work out their own types of convergent phenomenology, each suited to the needs of their respective fields. Perhaps "psychopsychology" could be used to refer to the convergent method of phenomenology applied to psychology.

I am afraid I must immodestly affirm Gopnik's remark that the pivotal ideas in the paper are mine, and not James'. Beyond drawing out some of the implications of the fringe for current research, I have tried to augment James' account with a functional analysis (section 4). Even more immodestly, I will assert that isolating and perusing a functional analysis of the fringe may have more long-term value than any of my specific findings, since I offer a concrete example of how to go about *explaining* phenomenology in functional terms.

Note that I use "function" and its cognates in their normal biological sense; this usage is altogether different from the "functionalism" now infesting philosophy of mind and cognitive science. Gopnik is quite right to say that my treatment of consciousness "probably will not allow us to reduce our experience to our functional states, it will not explain away consciousness."

One of the grave problems with contemporary functionalism is this tendency to reduce consciousness or, better, to *abstract* consciousness so that all that remains are certain "functional" relations, i.e., relations that can be instantiated in otherwise completely distinct information bearing media. But functional identity is not ontological identity. Questions about consciousness are simply not limited to questions about the information, expressed abstractly, that consciousness bears. In part, understanding consciousness means finding the particular *way* consciousness bears its information, just as part of understanding biological

Reply to Commentaries on B. Mangan (1993). Taking phenomenology seriously: The "fringe" and its implications for cognitive research. *Consciousness and Cognition*, 2(2), 89-108.

reproduction means understanding the particular way DNA bears its genetic information. As functionalism uses the term, DNA is "functionally" equivalent to all other media that can encode the same information, e.g., to an electronic pattern in a computer. But this does not change the fact that the aim of biological research is to understand the full ontological reality of DNA and the *specific* way DNA does its job. In other words, for functionalism, abstracting to the "functional" level necessarily obliterates specific aspects of biological ontology and thus denies the possibility of any scientific research on them.

On my account, then, the specific way consciousness bears context information is as vague experience, and the specific way consciousness executes a call function is by attempting to focus on its fringe aspect. We can imagine computers executing precisely the same call and context representation functions—but without using anything like this particular representational and command strategy to do it.

This brings me to Gopnik's one reservation about my treatment of consciousness. "There seems to be a confusion between the phenomenology and the underlying cognitive states. Mangan is perfectly right that evidence from both sources is germane to psychological investigation and that both areas may illuminate one another. Nevertheless, he often talks as if consciousness, itself, the set of subjective experiences with a particular phenomenology, were a cognitive system in its own right."

Guilty as charged. On my account, consciousness *is* a cognitive system, embedded in, and interacting with, a host of other nonconscious cognitive systems. And I claim that looking carefully at the fringe/focus structure of consciousness provides further reason to believe that consciousness does indeed have causal efficacy. For this reason, I would resist aspects of Gopnik's reformulation: "Mangan talks as if the cognitive systems were to be identified with the phenomenology. The fringe just does call information into consciousness. It seems to me that a better way of describing the relation would be to say that a cognitive system that calls unattended information into attention causes the phenomenology of the fringe." I presume that all our phenomenology is a manifestation of the brain and that the efficacy of consciousness depends on the brain. A full account of an act of attention calling new information into consciousness can never ignore the supporting, nonconscious neural systems. But at the same time, I would hold that a full account of successful, willed recall must, in general, also refer to the operations of consciousness, *qua* consciousness, as one part of the causal story. But let me postpone a more detailed discussion of the cognitive efficacy of consciousness until I consider Velmans' comments.

Gopnik brings out a point in my paper that, as written, is relatively minor, but should still interest any curious person: the phenomenological link between aesthetic and religious experience and its relation to normal cognitive activity. On my account, this link is "rightness," and its mundane and exalted forms differ only in intensity. This is not a claim James ever made, so far as I know. Early in life James toyed with becoming a painter. Yet oddly enough he almost never discussed aesthetic experience in his books and papers, nor did he apply his fringe analysis in the *Varieties of Religious Experience*. But then James detested

system builders: as Galin notes (1993), James did not even use the fringe in his theory of emotions.

But extending this discussion would take us far afield, and the reader can always look at Mangan (1991). I will only add that so far as I can see, my account does not *mandate* any position about the “ultimate” ontology of rightness, nor of consciousness in general. Perhaps mystic experience is nothing but a freak “spike” of intensity in an otherwise mundane cognitive mechanism. Perhaps ordinary conscious experience rests on a divine ground that manifests itself only when it moves toward its “natural” high intensity level. Needless to say, the former interpretation best fits the current scientific mind-set. However, I can see no reason to rule out completely the second interpretation. It is conceivable that consciousness has been somehow captured by the brain and yoked for biological ends. In any case, whatever consciousness’ ultimate status (and the above possibilities are hardly exhaustive), I believe the crucial research point is that, whatever its origin, consciousness is amenable to functional, cognitive analysis.

Bernard Baars’ commentary unpacks some of the empirical implications of my argument very clearly, and at the same time sets up a fresh context for the reader to get a sense of the fringe and its cognitive functions. Frankly, it is difficult for me to find more than a few quibbles with what Baars has to say. Point (a) of his operational definition of vague experiences requires that “all” subjects be able to report them. This seems too strict. “Most” would be more than adequate. We accept many psychological phenomena when evidence for them simply rises consistently above chance. Every subject in an experimental group does not need to show the effect for us to have confidence that the effect is real. Except as a politic compromise with prejudice, there is no ground for holding experiments on consciousness to a higher standard than that used in any other branch of related empirical research.

Another quibble: Baars says in point (v) of his discussion of a TOT that it “does not have . . . location [and] intensity.” But on my account, anyway, it is precisely the intensity of felt accessibility that serves as the datum for a FOK (feeling-of-knowing) confidence rating. The stronger the intensity of these feelings (in some cases rightness, in some cases familiarity, in some cases an overlay of the two), the more likely we feel we can retrieve the word in question. And when we strain to retrieve a word, I suspect we *do* have a sense of a roughly peripheral target–location relative to our body image. In my case, the feeling of grasping for a TOT often seems to be in a region above and behind my right ear and just outside my skull. (For a more general discussion of the phenomenology of the “backward retraction” of attention, see James’ remarks, including a long passage from Fechner, in the *Principles*, I, pp. 435–436.) Baars does say later that his point (v) should be qualified and ends his first section by largely retracting it.

I certainly agree with Baars that “*All* of our conscious experience, after all, is guided by . . .” the sort of mechanisms implied by TOTs. As I tried to show, TOTs are just the tip of the cognitive iceberg. Contrary to Baars, I think this was also James’ view. The aim of James’ discussion of TOT was to uncover general features of fringe and transitive experience; and on James’ model, experience always has a vague component, even during its “substantive” stage. James at

one point insists, for example, that “the image per se, the nucleus, is functionally the least important part of the thought” (p. 472).

In my paper I emphasized the rightness aspect of the fringe, and it is useful that Baars has begun to look at wrongness. When thinking about the fringe in operation, it is best to view it systemically, rather than just looking at a single, isolated experience. It is also important to keep in mind the fundamental distinction between the summary and call functions of the fringe, since being able to perform one function does not guarantee the successful operation of the other. It is often true that simply feeling that something is wrong (a summary evaluation) will not guarantee quick retrieval (a call) of the corrective information via the fringe. Of course, this is true for rightness, too: we can feel strongly that something is right without being able to say just what it is that makes it right.

Wrongness often initiates problem-solving. In the case of Baars’ example “The ship sailed past the harbor sank,” wrongness functions as the summary signal informing consciousness that tacit context conditions are not satisfied; but this signal is not in itself sufficient for us to go on and retrieve the specific grounds for the negative evaluation, and a good deal of conscious work in the focus of attention will still need to be done in such cases. But in many other cases the feeling of wrongness *will* lead us quickly to the specific source of the problem. In proofreading, for example, we often immediately feel that a word is wrong (misspelled) before we know precisely *what* the mistake is. And very quickly we retrieve a specific candidate that will hopefully correct the mistake. The retrieved candidate then gets a *subsequent* fringe evaluation; for in general, *after retrieval of a potential solution into consciousness, further global evaluation occurs*. (This implies that the processes responsible for preliminary global evaluations are rather coarse.) If a possible spelling feels strongly right once it actually enters focal attention, we make the correction without further ado; we do not unpack more detailed information to check this global evaluation. But if the candidate spelling feels strongly wrong when called to the nucleus, we reject it immediately and again start to grasp at the fringe for other possibilities. Then, too, there are the cases when we do *not* have a strong, unambiguous evaluative signal in the fringe. Here we fall back on a sequential process of detailed conscious examination and, perhaps, pick up our dictionary. This is a simple example of the sort of systemic approach to fringe experience I have in mind. Looking at a single experience at a single moment in time is rarely sufficient.

Baars wonders how well abstract concepts fit the fringe/focus model, since he takes abstractions to be at once focal *and* vague. As a logical matter, the situation is complex in detail: *to say that peripheral experience is vague does not imply that all vague experience is peripheral*, nor that because some vague experiences combine call and summary functions that all vague experiences do so, nor that those experiences which normally call information into consciousness will always succeed. In general, the contents of focal attention are clearer than *the fringe*. But this hardly means that the contents of focal attention have at all times a fixed degree of clarity—obviously there is substantial variation. Sometimes focal attention collapses altogether, and consciousness manifests a kind of mind-fog (see *Principles*, pp. 403–404). Baars is quite right (as is Galin) to point out that

at times the nucleus can be relatively ill-defined and more fringe-like in its phenomenology than at other times. Baars speculates along lines not unlike my own that in the vague phenomenology of abstract thought, there may be a trade-off benefit: for if the focus were more concrete, it would “take up more of the thinker’s limited [conscious] capacity;” the vague aspect Baars thinks especially prominent in our use of abstract concepts would then be freed for a richer representation of relations.

Certainly most cases of abstract thought would seem to have both fringe and focal components. And if we presume that abstract concepts have an especially large number of relations, we would expect the “fringy,” transitive character of consciousness to be relatively strong during abstract cognition. On the other hand, any concept, however abstract, that is designated by a name or other symbol has a clear-cut nucleus (Baars’ “handle”) around which the other vague and “relational” elements of the concept can cluster. Note that in the Einstein passage used by Baars, nucleus-like entities—substantive and articulated points of “arrival” as James would call them—are still evident: “the physical entities which seem to serve as elements in thought are certain signs and more or less clear images, which can be ‘voluntarily’ reproduced or combined.” So while Einstein takes these “signs and more or less clear images,” to be *idiosyncratic*, his cognitive phenomenology, as described, is well within the bounds of the basic nucleus/fringe model.

Max Velmans’ commentary raises a number of interesting questions. As he notes, it is based on an antecedent theoretical position: If we are to talk reasonably about consciousness Velmans believes we must distinguish three different senses: (1) the actual subjective feelings of which we are conscious—in other words, to our phenomenology; (2) nonconscious mechanisms whose “products are conscious (Velmans’ emphasis)—that is, processes which produce our phenomenology; and (3) any cognitive process in which consciousness is presumed to be “causally efficacious.” Velmans believes he has shown (Velmans, 1991a,b) that “from the perspective of an experimenter or external observer, no human information processing is conscious in sense 3.” If true, this is a monumental finding. We will consider the grounds for it below.

Velmans, then, begins with the premise that I must be wrong about one of my major conclusions: that the fringe provides a new piece of evidence against epiphenomenalism. But otherwise we have few points of conflict. So, for example, Velmans has no objection to my connectionist interpretation of rightness, except for its interactionist implications. In fact, he begins his commentary by giving a good sketch of the background experiences of rightness and wrongness, of understanding and not understanding, etc. and notes that a “simple way to make sense of the ‘functional significance’ of such feelings is just the way Mangan suggests,” that is, by explaining these feelings as global signals of context fit or mismatch.

Velmans is even willing to go further in my direction: “It may be . . . that fringe feelings of ‘rightness,’ ‘understanding,’ and so on *represent* global evaluations of the state of the cognitive system (generated by neural activity), and that such evaluations guide attentional processing in just the way Mangan suggests. If so,

models of attentional processing need to include such global evaluations, along with the mechanisms for generating them. The global evaluation/focal attention distinction may also help to explain cases where these appear to be dissociated, in blindsight, anosognosia, and so on" (Velmans' emphasis). But even if all of this were accepted, Velmans adds, "one needs no commitment to the view that the fringe feelings actually *cause* brain processing to operate in ways that cannot be understood in purely neural [nonconscious] terms" (Velmans' emphasis).

This is quite true. We must be very clear about the logic of the situation. My analysis does not claim to *prove* that consciousness has cognitive efficacy. In general, just because we feel that conscious feelings have a real effect on our subsequent behavior and thoughts, we cannot *necessarily* conclude that these feelings are in fact doing what they seem to be doing. Other interpretations are possible.

But we are rarely concerned with the mere logical *possibility* of a theory in science. (This is one major difference between science and philosophy.) In science, theories of significant scope are almost always less than perfect, and in any case one can almost always construct an alternative theory to account for a given set of findings. But various theories *do* differ as to their plausibility, relative to the contending alternatives. As new evidence and other considerations arise, relative plausibilities can shift to such a degree that for a time, anyway, one theory is clearly to be preferred over all the others. In other cases, it is difficult to decide among the contenders.

Now I *do* claim that by examining the systemic interplay of focus and fringe, we find that consciousness contains a set of experiences that certainly *look* like they are designed to mediate conscious/nonconscious interaction. Since far and away the most plausible account of nonconscious processing is neurobiological, fringe analysis then gives us new reason to believe that conscious/neural interaction occurs. Do I think this line of reasoning conclusive? No. But my analysis does strengthen the view that consciousness has causal efficacy and weakens the opposite view. Of course, this presumes that the cognitive efficacy of consciousness is still an open question. Certainly, in my opinion, Velmans' arguments against the cognitive efficacy of consciousness are open to considerable dispute.

This is not the place to review my critique of Velmans (Mangan, 1991) in great detail. Let me just note that Velmans' general approach (1991a) is to bring out cases where cognitive functioning occurs when we might otherwise assume it to be mediated by consciousness, but for one reason or another the functioning seems in fact to be executed in the absence of conscious experience. There are two major difficulties here: (1) Biological systems often overlap and provide an organism with more than one way to execute a given function. The capacity to execute a function in the absence of consciousness in some unusual (and fairly primitive) situation does not show that, in normal circumstances, consciousness' cognitive role is different from what it appears to be. (2) This consideration is strengthened when we look at the actual kinds of functioning/awareness dissociations Velmans is able to point to. They are drawn from relatively low-level sorts of processing, such as the subliminal interpretation of a well-known word or, at

best, a simple sentence. But it has been noted often that consciousness seems to be most typically active in novel circumstances (Baars, 1988), and examples of robust novelty processing are conspicuous by their absence in Velmans' treatment. In addition, if fringe feelings like rightness do occur, many of the experiments cited by Velmans as instances of a functional/awareness dissociation (e.g., blindsight, subliminal perception) could well turn out to involve conscious experiences after all.

Then there are the far more general "philosophical" considerations. Velmans remarks that I do "not specify *how* fringe or focal consciousness might perform" their functions in the sense of being able to causally influence neurons. "As he [Mangan] admits elsewhere, 'There is still the mystery of the interaction of consciousness with the physical brain' (Mangan, 1991)." Again, Velmans is perfectly right. I openly admit that an interactionist account (which is not necessarily "dualist," by the way, although this is another question that would take too much space to pursue) involves a mystery. It would be folly to pretend otherwise. The problem is that all the *other* positions involve mysteries of one kind or another, too. The best scientific approach at this point in cognitive research is to look for the position that minimizes the mystery, does the least violence to the general canons of scientific thinking, and maximizes the scope of what we *can* explain.

With these desiderata in mind, I would point out that once we admit that there *is* a distinct cognitive realm of subjective experience, we do the least violence to our scientific understanding of the human organism by presuming that the brain *somehow* produces consciousness. The "somehow" is the mystery, but it is a mystery that even epiphenomenalists must live with, since they must still explain where consciousness comes from. They, then, already admit the primal mystery: the brain causes consciousness. But once this is posited, denying consciousness a cognitive function becomes perverse, since consciousness manifestly seems, from the inside, to execute cognitive functions, and from the outside it is a presumptive biological product: the apparent "subjective" functionally is just what we would expect of consciousness as an "objective" outcome of evolution. The more system-like consciousness seems to be, the less likely it can be dismissed as a functionless mutation, and the more it looks like it is shaped by a long selective process. But being subject to an evolutionary process implies that consciousness is functional: it offers some adaptive advantage, otherwise evolution could not shape it. The standard epiphenomenalist, then, pays the basic tax of admitting a mysterious connection between consciousness and brain-matter, but without getting any explanatory, biological return. For the epiphenomenalist will not pay the relatively slight additional cost of presuming that the brain that produces consciousness can in turn be influenced by consciousness. Epiphenomenalism thus cuts itself off from the capacity of using functional, biological reasoning to investigate consciousness. Relative, then, to epiphenomenalism, an "interactive-functional" stance of the sort I am taking toward consciousness is clearly a better candidate for satisfying the scientific desiderata.

Velmans denies he is an epiphenomenalist. "I do *not* adopt epiphenomenalism," Velmans says, "for the reason that it has nothing to say about causality

viewed from a first-person perspective'' (1991a). I must confess I am a bit unclear about Velmans' general relation to epiphenomenalism. Certainly he holds one defining tenet: consciousness does not influence the brain. But he sees that the standard epiphenomenalist view entails admitting *some* relation between the brain and consciousness, and even aspects of this link seem to bother him: "Indeed, it is not obvious how events as apparently different as brain states and experience *could* have causal influence on one another. As Mangan points out, this problem applies to epiphenomenalism as much as to dualist–interactionism" (Velmans, 1991a). Apparently, Velmans wants to remove consciousness as far as possible from any sort of adaptive, functional analysis. Significantly, this means locating consciousness even further from the domain of biological explanation. Velmans does not want to accept "that the only function or goal worthy of consideration is survival fitness. To avoid the impasse one has to re-examine the presuppositions. Viewed from a third-person perspective, consciousness does not enhance adaptive functioning. Rather, the brain functions, in part, to produce experience. From a first-person perspective, the difference this makes is obvious. Without consciousness there would be no experienced world. . . . From this perspective, consciousness *is* the goal—but not one that can be understood in (third-person) information processing terms" (Velmans, 1991a; Velmans' emphasis).

It would seem, then, that the mystery embedded in Velmans' view of consciousness is at least as deep as that found in a scientific form of interactionism (i.e., that consciousness results from a *physical* brain process we have yet to discover, and that consciousness *physically* influences the brain that manifests it). On Velmans' account, to remove this possibility, we must give up standard modes of biological explanation when it comes to first-person experience. This move then produces substantial and mysterious "lacunae" of its own. As we saw above, Velmans denies that consciousness is subject to the adaptive, selective mechanisms that presumably shape all other aspects of our organism and keep them working together. This means that he puts an *additional* mystery into the heart of his position, quite beyond the assumption that the brain produces consciousness. For if consciousness is in fact really distinct from all attentive processing mechanisms, what keeps consciousness so closely linked to attention? Velmans can only respond: "The conditions for consciousness in the human brain are (in part) *produced* by focal-attentive processing. That is why they correlate so closely" (Velmans, 1991a; Velmans' emphasis). This means that for some mysterious reason, a very small subset of neural activity (that responsible for focal attention) *just happens* to produce consciousness. For Velmans, no third-person account of biological, evolutionary advantage can *in principle* explain how the consciousness/attentive process correlation is maintained, since there is no way selective pressure would be able to keep them closely correlated.

The lacunae in a physical interactionist view of consciousness do far less violence to scientific ways of thinking than either Velmans' position or that of more orthodox epiphenomenalists. For the physical interactionist, our inability to explain how the brain produces consciousness and how consciousness influences the brain are nothing but simple, honest ignorance. Given the history of science,

we have good inductive evidence that the mystery of consciousness will, in time, yield to the advances of empirical research. It is important to see that far and away the most basic and (currently) inexplicable mystery—that the brain somehow produces consciousness—is also accepted by Velmans and the epiphenomenalists. All three pay the same basic mystery tax. But interactionism goes on to accept the possibility of a symmetrical, feedback relation: that consciousness in turn effects neurons. Velmans and the epiphenomenalists deny this symmetrical possibility. Again, remember that Velmans himself is clear that from a first-person viewpoint, interaction *seems* to be the case, i.e., we *seem* to translate wilful conscious decisions into what science tells us are physical, neurally mediated actions.

The only *special* tax exacted by physical interactionism is the admission that current science is not omniscient and still has much to learn. The distinguishing “mystery” of interactionism goes no further than the fact that science, at this moment, has less than complete knowledge of the brain. Benefits of holding physical interactionism include the fact that it can easily incorporate more first-person evidence than the contending theories (e.g., that consciousness feels like it has cognitive efficacy), and that it can naturally utilize existing and successful modes of biological investigation. Again, the general, and by far the largest, mystery tax is paid equally by Velmans, orthodox epiphenomenalists, and physical interactionists alike: they all accept, without benefit of a specific explanation, that the brain produces consciousness. But epiphenomenalism gets no explanatory or other research utility for *its* payment. It is inherently a negative and sterile position. Indeed, if accepted, *epiphenomenalism would stop all research aimed at exploring interactionism*. Any position that a priori excludes an entire domain of research possibilities is obviously a danger for scientific thinking and should be held warily, as a last resort. While Velmans’ position is in some ways far more sensitive to consciousness than standard epiphenomenalism (for he takes first-person experience in itself very seriously), in this context his view is perhaps even more pernicious. For insofar as I understand his position, Velmans denies the possibility of *any* full, scientific, third-person explanation of first-person experience. This would have us give up biological modes of consciousness investigation de jure, as it were, something quite beyond the simple de facto dismissal of consciousness research by less sophisticated epiphenomenalists. (However, it must not be forgotten that Velmans’ asymmetrical stance does let him use first-person knowledge to augment third-person knowledge. In balance, then, I think it clear that physical interactionism best satisfies the scientific desiderata.

David Galin argues that the terminology drawn from James needs to be abandoned. Galin offers alternatives that he feels are more “congenial to cognitive psychology” and that avoid importing possible phenomenological errors and other misconceptions into my analysis, especially problems around the notion of vague experience. This may have its valuable side. We certainly do not need to accept James’ terms as final and should always be open to the possibility of refining a phenomenological analysis and working out a better nomenclature. Although James is one of the finest stylists to write on consciousness in English,

there is always room for improvement. In one sense, this is just what I have tried to do with my functional explanation and extension of James' model. Perhaps, as James' phenomenological points are grasped through the medium of his vivid and evocative prose, we should have a more technical vocabulary ready in the wings to help elaborate distinctions, bring out points that have been, or risk being, confounded, and so on.

And I believe Galin makes a useful point about the need to clarify the apparent "dichotomy" implied by the nucleus/fringe distinction—although he takes these terms to imply a stronger split than either James or I intend. Since Mangan (1991) I have used the notion of consciousness' limited "articulation" capacity to help bring out the unifying trade-off relation between nucleus and fringe. But I have probably applied this point more conspicuously to functional questions than to the phenomenology itself, and Galin's remarks make clear that need to discuss more fully the phenomenological implications of consciousness' limited articulation or, alternatively, "resolution" capacity. (To do this adequately would require much more space than is available here, but I will try to give some sense of my view of the relation of vague to clear experience shortly.)

Unfortunately, at a more specific level, Galin's proposals generally seem to me either unnecessary or wrong. Let me consider first narrower questions of terminology and then the more substantive problems about how to characterize vague experience.

Galin would change a number of terms that, in my view, are already adequate. His new terms seem to me to offer little or no improvement and are often problematic. The result of such a wholesale switch would do little more than obscure the actual origin of these distinctions; I doubt they would give us any real clarity or new insight into the entities Galin would rename. So, for example, Galin proposes that "The term 'fringe' should be replaced." His candidate is "non-feature awareness." It is difficult to see the gain. Besides being awkward, as Galin notes himself, the term would create more problems than it solves, especially given the current connotations of "feature" in cognitive psychology. Features are often said to be *integrated* into objects in focal attention ("focal attention" is almost a precise synonym for James' "nucleus"), but individual features are said to reside in preattention. We saw in my theoretical focus article that the preattentive realm is ambiguous as to its conscious status, but in either case, Galin's "non-feature awareness" would be an oxymoron, since one of the main functions of "non-feature awareness" would be to signal the existence and accessibility of features! Recall my discussion of Rock's experiments. "Inattentive" experience has features, but they do not enjoy formal integration. If "fringe" is found to be an inadequate term, the better choice would be inattention, since it is already used in the cognitive literature to refer exclusively to conscious awareness that is not attentive, and on Rock's analysis, it further involves the notion of potential information that can be called into attention. Another alternative is "non-focal" awareness. But I can see little force in Galin's argument that any change from "fringe" is needed.

Some of Galin's other proposals are very slight reworkings of James' terms

and, although they raise no confusions, seem unnecessary. So, for example, Galin proposes "relational awareness" for James "feelings of relation." In the same sentence, Galin also proposes the term "summary awareness," which I believe could be troublesome, since it risks making just the sort of confound Galin is trying to avoid; in this case, confounding the conclusions based on a functional analysis of the phenomenology with the phenomenology itself.

Galín notes that calling fringe states "vague" might imply to the unwary that these feelings are only felt in weakened form, when, on my account at least, vague experience at times can be extremely intense (e.g., during aesthetic experience). Here I think Galín is on firmer ground. Without further specification, one might well assume that a vague or fringe experience *would* have relatively low intensity. But simply telling the reader directly that the fringe is important and that a vague experience can be intense without being clear is, I think, enough to override any further misunderstanding. *Intensity is not articulation.*

To justify many of his proposed changes, Galín asserts that "Because Mangan is concerned primarily with the feeling of rightness, he . . . does not emphasize any further distinctions. Neither James nor Mangan attempted to sort the very heterogeneous array of fringe awareness into categories, either phenomenologically or functionally. Thus they implicitly gave us the dichotimization of awareness into nuclear vs everything else." Frankly, this justification for changing the terminology seems to me as odd as the proposed changes. Even in the relatively narrow compass of my position paper, I believe I have developed (beyond rightness) the notion of control experiences in the fringe, explained the summary vs. call distinction in fringe experience, and explicitly contrasted the phenomenological description of the fringe (section 3) with its functional analysis (section 4). Nor is James' treatment of the fringe the mishmash that Galín suggests. Readers are encouraged to consult James, especially the "Stream of Thought" chapter, for themselves.

Galín does recognize that more needs to be done on the nucleus/fringe distinction. Unfortunately, Galín's specific attempts in this direction also seem to me to be mistaken and misleading. Aside from the more narrow terminological problems, Galín confounds linguistic and logical features with phenomenological features (e.g., in his example of expressing distance variously between Paris and goalposts) when discussing "vague" experience. It is certainly true that both a verbal description and an experience can be *said* to be vague or clear based on one's purposes. And some purposes demand greater phenomenological detail than others. But the *phenomenological* point is that, whatever the purposes, experience is almost never a homogenous field. It almost always has a region of focus (today called focal attention) and there does seem to be a region of less "articulated" experience that in a spatial sense does seem to surround it. Of course this structural feature of consciousness is quite plastic, as noted in the earlier discussion of abstract experiences in my remarks on Baars' commentary.

Neither James nor I treat the nucleus and fringe as altogether distinct entities: they do have different functional roles in cognition and typically have distinguishable phenomenologies, but they still blend and intermingle with one another, on

both our accounts, into one continuous entity. In James' stream of consciousness metaphor, the pots and buckets and the rest (i.e., the specific contents) were themselves immersed in the "free water" of vague feelings that surround them, and *filled them as well*. James often talks about the unity of consciousness and insists that the fringe aspect is "fused into one with it [the nucleus] and has become bone of its bone and flesh of its flesh" (p. 255).

On my account, the biological analogue to the allocation of pixels is the result of a complex trade-off strategy: some pixels go to the menu-bar and, among other things, mediate voluntary retrieval; but the bulk of pixels are used to "articulate" the immediate task at hand. Applied to James' dynamic transitive/substantive model, this means that the "pixel" allocation in consciousness between vague and clear experience is constantly changing. Sometimes more pixels work to represent relations, so that fewer pixels are available to articulate the entities in focal attention, and this relation is in constant flux. Furthermore, consciousness' "pixels" seem to come in different sizes—moving by degrees from a fine-grained structure at the focus of attention to a more and more coarse grained or "fat" pixel structure toward the periphery. Thus "vague" and "clear" experiences are in one sense distinct regions of experience and yet merge, at the phenomenological level, into a unified field created by a limited by articulation capacity mechanism. Among other features of the analogy, we can see how an experience can be vague but intense.

Katharine McGovern's commentary shows the wide range of potential applications of the fringe model of consciousness. It suggests that more explicit phenomenological research will, perhaps, be useful for the study of emotion and related phenomena. Her review touches on a wide domain of cognitive phenomena and it is of course suggestive that she has been able to find a unified strand running through so many areas. I would caution, however, against using "featureless" to characterize the fringe aspect, since at the least this risks creating a confound between usage in the study of emotion and established usage in more "hard line" cognitive psychology.

As McGovern sees, one value of taking phenomenology seriously is that it may let us bridge the gap between the ordinary modes of expression used naturally by people in both interview and therapeutic discourse. It seems possible to integrate "feeling" naturally into a reasonably detailed cognitive model, and more generally bring the most subjective aspects of human experience within the reach of a linkage with laboratory findings. It would be interesting to see if research into the use of emotive language and laboratory findings could let us work out a map of the strata of increasing articulability, and their typical contents.

McGovern's report of the Greenberg and Safran (1987) study is especially welcome. It is fascinating to see the independent link reported between subjectively felt meaning and TOTs. As McGovern notes, by understanding the phenomenologically triggered retrieval mechanisms, it may be possible to develop new therapeutic techniques. Work at a more "inarticulate" but primary evaluative level of personal exploration might proceed better when unavailable material is allowed to enter and stay in consciousness *without* aiming for immediate retrieval into

consciousness. Just as we can learn to retrieve a portion of a memory through focused attention, “defocused” attention may be helpful in allowing otherwise unavailable material to stay available but in an inarticulate fringe state.

REFERENCES

- Baars, B. J. (1993). Putting the focus on the fringe: Three empirical cases. *Consciousness and Cognition*, 2(2), 126–136.
- Baars, B. J. (1988). *A cognitive theory of consciousness*. London/New York: Cambridge Univ. Press.
- Galín, D. (1993). Beyond the fringe. *Consciousness and Cognition*, 2(2), 113–118.
- Galín, D. (1993). *The structure of awareness: Contemporary applications of William James' forgotten concept of the fringe*. Submitted for publication.
- Gopnik, A. (1993). The psychopsychology of the fringe. *Consciousness and Cognition*, 2(2), 109–112.
- Greenburg L. S., & Safran, J. D. (1987). *Emotion in psychotherapy*. New York: Guilford.
- James, W. (1890). *The principles of psychology*. New York: Dover. (Reprinted in 1950)
- Mangan, B. B. (1991a). *Meaning and the structure of consciousness: An essay in psychoaesthetics*. Doctoral dissertation, University of California, Berkeley. (University Microfilm No. 92033636)
- Mangan, B. B. (1991b). Epi-arguments for epiphenomenalism. *Behavioral and Brain Sciences*, 14, 689–690.
- McGovern, K. (1993). Feelings in the fringe. *Consciousness and Cognition*, 2(2), 119–125.
- Velmans, M. (1993). A view of consciousness from the fringe. *Consciousness and Cognition*, 2(2), 137–141.
- Velmans, M. (1991a). Is human information processing conscious? *Behavioral and Brain Sciences*, 14, 651–669.
- Velmans, M. (1991b) Consciousness from a first-person perspective. *Behavioral and Brain Sciences*, 14, 702–726.