

Review

Reviewed Work(s): Consciousness Explained by Daniel C. Dennett

Review by: Sonia Sedivy

Source: Canadian Journal of Philosophy, Sep., 1995, Vol. 25, No. 3 (Sep., 1995), pp. 455-

483

Published by: Cambridge University Press

Stable URL: https://www.jstor.org/stable/40231921

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at https://about.jstor.org/terms



Cambridge University Press is collaborating with JSTOR to digitize, preserve and extend access to Canadian  $Journal\ of\ Philosophy$ 

Critical Notice

DANIEL C. DENNETT, *Consciousness Explained*. Boston: Little, Brown 1991. Pp. vii + 511.

# Consciousness Explained: Ignoring Ryle and Co.

If my argument is successful, there will follow some interesting consequences. First, the hallowed contrast between Mind and Matter will be dissipated, but dissipated not by either of the equally hallowed absorptions of Mind by Matter or of Matter by Mind, but in a quite different way.

... both Idealism and Materialism are answers to an improper question. The "reduction" of the material world to mental states and processes, as well as the "reduction" of mental states and processes to physical states and processes, presupposes the legitimacy of the disjunction "Either there exist minds or there exist bodies (but not both)." (Gilbert Ryle, *The Concept of Mind*, 22)<sup>1</sup>

Consciousness Explained is Daniel Dennett's explicitly avowed attempt to provide a 'conservative materialistic theory' (140) of conscious mind. It is an attempt driven by the disjunction that 'Either there exist minds or there exist bodies (but not both)'; or, in Dennett's words, that 'Once we take a serious look backstage, we discover that we didn't actually see what we thought we saw on stage' (434). Thus Dennett's project is not simply the constructive one of showing that 'the brain must be the mind' (41) and 'how the mind is accomplished by the brain' (322). The project is at once also explicitly and avowedly destructive.

I will argue that this is not a position that Dennett, given the theoretical commitments driving much of his previous work, should be elaborating. Moreover, aside from the commitments of his own previous work, the

<sup>1</sup> Chicago: University of Chicago Press 1949

sort of theory Dennett advances in *Consciousness Explained* attempts to answer a misconceived demand, the demand to supply a *reductive*, materialist explanation of consciousness. However, the Pandemonium Multiple Drafts model of brain processes which Dennett develops need not be used for reductive purposes. The empirical merits of a model of brain functioning do not speak to the issue of reductionism, and so cannot be used to recommend a reductionist framework in which to locate the model. Hence, because my objective is to assess the theoretical framework in whose service Dennett advances the Multiple Drafts model, I will not address the empirical adequacy or advantages of the model. Insofar as the Multiple Drafts Model provides excellent resources for explaining what is happening backstage, as many of us suspect or believe, it nonetheless would not require us to alter our appreciation of the happenings onstage.

Dennett's theatrical metaphor stresses that the differences and relationships between what happens on stage and off don't need to be explained in terms of metaphysically different kinds of stuff. In this way it works against him, highlighting the fact that the argumentative strategy of the book is misdirected. Dennett chooses to target Dualists and Cartesian Materialists who believe that the demand for a *reductively* materialist theory should be answered differently from the way he proposes, and to ignore those who believe that the demand is misconceived and so is not to be answered at all. This choice is strategic in that it allows Dennett to avoid discussing certain difficult issues in a work ostensibly written for nonprofessionals. However, I believe that the strategy backfires.

By the 1990s, neither Dualists nor Cartesian Materialists make up the most philosophically significant opposition to Dennett's reductively materialist theory. Rather, the opponents, amongst whom I always used to place Dennett, are those who believe that understanding ourselves as experiencing subjects is precisely that, a distinctive form of understanding or explanation which cannot be mapped onto another distinctive form of understanding of ourselves as made of physical processes governed by general laws. Today the duality at issue is not that of stuffs but of kinds of explanation. The Dennettian distinctions of personal and subpersonal levels, of semantic and syntactic engines, and of intentional, design, and physical systems and stances are sophisticated conceptual tools for dealing with the fact that certain physical things, namely ourselves at least, need to be understood and explained in more than one distinctive way. These concepts and distinctions go missing in Consciousness Explained. They are absent not only terminologically, owing to the nontechnical nature of the book, but they are missing theoretically in that Dennett's theses about consciousness fail to respect and even conflict with them. Consciousness Explained attempts to use subpersonal facts

about the design system to change what we countenance as the facts of conscious experience at the manifest personal level.

What is really hard in the late twentieth century is not, as Dennett's choice of focus suggests, to let go of our Dualist or Cartesian Materialist inclinations. Rather, it is to let go of the hankering to build bridges, to leave explanations unaligned and to stop feeling that there is some explanation that we are failing to supply when we countenance the requisite plurality of distinct explanations of ourselves. What is hard is to countenance fully the notion of distinctive forms of intelligibility and of facts which require being made intelligible in distinctive ways. Indeed, to understand that — one and the same thing — metaphysically speaking, may require being made sense of in more than one distinctive way.

But this is the challenge set for us by mid-twentieth century diagnoses of the spectre of a mind/body problem awaiting solution, diagnoses offered principally by Ryle, Wittgenstein, and Sellars.<sup>2</sup> While we encounter similar challenges in numerous domains as the developing scientific image becomes ever more distinct from the one which is manifest, the challenge seems especially difficult with respect to ourselves. This is because, of course, the manifest image of ourselves doesn't simply posit conscious experience but is itself consciously experienced. We consciously experience ourselves as consciously experiencing. But how could this phenomenon be akin to an explanatory posit? In offering a reductive or 'conservatively materialistic' answer to this puzzle, Consciousness Explained departs from the diagnoses of the mind/body problem offered by Ryle, Wittgenstein, and Sellars with which Dennett casts it as continuous.

To sum up these introductory remarks with more theoretical precision, I am claiming that dualism and metaphysical reduction, the reduction of mind to matter, are not the pressing theoretical issues today. The issue is explanatory and ontological reduction: the reduction of one type of explanation to another and of the entities posited in one type of explanation to the entities posited in another. The prevailing framework within which most theorists work countenances only one kind of stuff, namely physical stuff. Consciousness Explained miscasts the debate as one concerning metaphysical reduction and proceeds to reduce explanations and entities in the guise of metaphysical parsimony. That is, Dennett argues for identifying conscious mental episodes with certain functional

<sup>2</sup> Gilbert Ryle, The Concept of Mind; Wilfrid Sellars, 'Empiricism and the Philosophy of Mind,' Science, Perception and Reality (London: Routledge & Kegan Paul 1963); Ludwig Wittgenstein, Philosophical Investigations, trans. G.E.M. Anscombe (Oxford: Blackwell 1953)

organizations of brain states. On the ground of this identity, Dennett argues that since brain states lack some of the features of the manifest image of conscious experience, the manifest image is incorrect and requires alteration. I will argue against the latter move by showing that the former is not established. Dennett does not provide an adequate case for the identification of conscious mental episodes with functional organizations of brain states and hence he lacks grounds for using facts about the brain's functional organization and activities to deny features or facts of manifest conscious experience. Dennett's case for the identification of conscious mental episodes with brain states is inadequate because it does not address considerations that suggest that mental episodes and brain states figure in distinctive explanations of ourselves.

My examination of Consciousness Explained falls into five parts. I begin with a summary of Dennett's constructive, empirical theory of mind. The second section outlines the destructive argumentative strategy of Consciousness Explained directed against much of our conception and understanding of consciousness. The third section summarizes Dennett's arguments for the reductive theoretical framework within which he needs to locate his constructive empirical theory in order to draw his destructive inferences about consciousness. Dennett's case for his reductive framework is critically examined in the fourth section. I focus on considerations for the plurality of distinctive explanations of ourselves, as they have been articulated in Dennett's earlier work, and argue that Dennett's reductive framework is not adequately supported in the face of these considerations. In the concluding fifth section, I examine Dennett's chief destructive inferences about conscious experience. Having argued that the reductive framework is not adequately supported, in this section I focus on the further substantive assumptions Dennett needs for these destructive inferences.

## I The Constructive Theory: The Multiple Drafts Model

Dennett begins by describing human conscious experience as a 'Joycean' stream of judgments. He goes on to explain that our streams of judgments are the states of a serial virtual machine and that, although the states of the virtual machine make up a sequence, they are in fact implemented by a Pandemonium of competing parallel brain processes. That is, a 'Joycean' stream of consciously experienced judgments *is* a sequence of states of a serial virtual machine that is 'run' or implemented in our biological hardware.

The computer analogy allows Dennett to apply the distinction between 'basic fixed or hard-wired architecture' and virtual architecture to

the brain and its workings. The brain is likened to a fixed or hard-wired computer architecture while the conscious mind is identified with the virtual machine that is functionally realized therein. The fixed architecture of the brain consists of many parallel processes acting at once. These parallel processes are organized into a virtual architecture or software that is serial in nature. Dennett's explanation *identifies* the sequence of states of this virtual machine with the sequence of judgments making up the stream of our experience.

According to Dennett, the function of many of the brain's parallel processes is that of 'content fixation.' Multitudes of parallel processes bear 'simple' contents from which ever more 'complex' coalitions bearing more 'complex' contents form. Ultimately, some of the 'complex' contentful coalitions make up the serial states of the serial 'virtual' machine. They do so by winning out in a free-wheeling undirected competition amongst the other 'complex' contents all being composed from alliances of 'simpler' contentful states in parallel. Since Dennett likens the multitudinous simpler contentful states to content 'demons,' the whole mess of coalitions of increasingly complex contentful states that are in competition amongst one another is a 'Pandemonium-style' model of content fixation. Dennett suggests that we understand the Pandemonium of content-fixations as Multiple Narrative Drafts from which a single, serial, continuously revised and perpetually revisable narrative is selected in undirected competition. Because of the undirected competitive nature of the selection process, the winning narrative is open to continual change or 're-editing' as different coalitions of content demons may become victorious. This winning narrative is the sequence of judgments which makes-up our experiential stream.

The fact that a single narrative draft is selected from multiple candidates — or in more technical terms, that a serial virtual machine is realized in multiple parallel processes — is explained in terms of a process of self-stimulation or probing. Dennett suggests that our brains probe themselves for information. These 'probings' function to select an 'answer,' thus yielding something like a question and answer sequence. Why would brains have a self-probing function? Dennett proposes that if brains have an overarching function, it is to 'produce future' by knowing what to think about next. That is, their function is to anticipate the future so that an organism's chances of success in its environmentare increased by allowing it to behave appropriately in anticipation of rather than simply in response to a course of events. To 'produce future' brains need to know what to think about next, which means that they need to be able to represent the relevant facts. According to Dennett, brains solve the 'meta-problem' of knowing what to think about next, of thinking about the right fact at the right time, by engaging in self-stimulation. That is, brains ask themselves questions

in order to produce answers. Thinking about the right thing is the result of answering one's own question.

Animals are capable of some degree of self-stimulation and in our case, that capacity exploded with the development of language and cultural products or 'memes.' Language developed from and in turn escalated our self-stimulatory activity into 'trains of thought.' In being taught a language we are taught habits of reasoning, of asking and answering questions, habits that become so ingrained as to become transparent to us as part of our nature. As a result, we are just about continuously engaged in asking ourselves questions and answering them, probing ourselves so habitually that we cease to be aware of doing so. Moreover, language is important in that in learning it '[w]e somehow *install* an already invented and largely "debugged" system of habits in the partly unstructured brain.'

We install an organized and partially pretested set of habits of mind ... in our brains in the course of early childhood development.... the overall structure of the new set of regularities ... is one of *serial chaining*, in which first one "thing" and then another "thing" takes place in (roughly) the same "place." This stream of events is entrained by a host of learned habits, of which talking-to-oneself is a prime example. (221)

So this, in a nutshell, is Dennett's account of conscious experience: the sequence of winning judgments in continuous self-probing that is implemented, in our case, in our brains' parallel fixed architecture.

Anyone or anything that has such a virtual machine as its control system is conscious in the fullest sense, and is conscious *because* it has such a virtual machine. (281)

### II The Destructive Argumentative Strategy

At the outset, I suggested that Dennett's aims are at once constructive and destructive. The above constructive theory is used to argue against much — if not all — of our theoretical as well as commonsense grasp of the nature of conscious experience:

- (i) Consciousness is neither the product nor the by-product of nonphysical stuff.
- (ii) Neither is it to be explained by Cartesian Materialism, which misconceives the way in which consciousness is realized in the brain. Cartesian Materialism identifies conscious states with spatially or functionally defined culminating states of brain processes. Thus Cartesian Materialism is the view common to much neuroscience 'that there is a crucial finish line or boundary somewhere in the brain, marking a place where the order of arrival

equals the order of "presentation" in experience because what happens there is what you are conscious of (107). Metaphorically speaking, Cartesian Materialism posits a theater where conscious states each take their turn. According to Dennett, it is the positing of the theater — of a process issuing in states that are fully determinate and conscious — which renders this form of straightforward reductive materialism Cartesian. What is Cartesian about this reductive materialism is that, though the theater and so all the goings on realized therein are made of bona fide (physical) stuff, what happens on-stage is conceived to be fully determinate events with fixed temporal and qualitative properties of which the viewing audience can therefore have an evident, certain grasp. It is Cartesian Materialism's identification of brain processes with episodes that are fully determinate (both temporally and qualitatively) to which Dennett objects.

(iii) Conscious experience is neither qualitatively nor temporally continuous and determinate; it does not consist of a continuous, fully determinate stream of qualitative states or episodes. Our manifest image of conscious experience as a stream of determinate episodes that have a continuous qualitative character and that occur continuously in a determinate temporal order, and so *any* theory premised on that image is misconceived as well. Not only are most of the phenomenal properties ascribed to conscious experience not the properties it really has — such as the property of determinate temporal order — but it is an error to ascribe phenomenal properties to conscious experience at all! There is no such thing as phenomenology, not really. Conscious states do not have qualitative character nor do they occur in a determinate temporal order: such properties are simply fictional.

I have suggested that there is no need to object to or to discuss Dennett's denial of dualism. Dennett's rejection of Cartesian Materialism brings together two different sets of objections: objections to a certain empirical theory of the brain's activities and objections to the commonsense image of mind which that empirical theory attempts to explain. Cartesian Materialism accepts the manifest image of mind and attempts to show how it might be realized in the brain. Dennett wants to deny the manifest image (iii) and so he denies the empirical theory which shows how the manifest image might be realized in the brain (ii). If the Cartesian Materialist *empirical* theory is incorrect then that set of Dennett's objections would be vindicated. However, in the absence of a reductive framework, evidence against the Cartesian Materialist empirical theory would not establish that consciousness does not have its manifest fea-

tures. Hence, insofar as Dennett wants to use empirical facts to show that certain aspects of the manifest image of consciousness do not obtain, he needs to establish a reductive framework. If his case for a reductive explanatory framework is not adequate, then the empirical work does not warrant denying what it is like to be consciously aware of one's world and that there is something which it is like to be consciously aware of one's world.

Dennett's specific arguments concerning conscious experience take the following form: facts about the subpersonal virtual machine are used to support inferences about which features may or may not be present at the personal conscious level. Hence the legitimacy of the specific inferences, for example about temporal or qualitative determinacy, depends in the first instance on the legitimacy of the type of inference which proceeds from subpersonal to personal facts. Accordingly, Dennett begins by presenting a general case for the legitimacy of (reductive) inferences from subpersonal to personal facts in Part I of the book, proceeding to the specific arguments in Parts II and III. So, let's follow his strategy, examining his general case and then proceeding to the specific arguments which utilize and rely on that theoretical framework.

#### III The Reductive Framework

Dennett presents a two-part case to legitimize using his subpersonal explanation to change our understanding of the personal level phenomenon of consciousness. In the first place, he argues that since dualism is just defeatism, the brain must be the mind. This would allow him to hold that features of mind are either identical to brain events or effects of brain events, so that discoveries about brain events show what properties the mind may and may not have. The second move is to cast our understanding of mind as *hetero*-phenomenological rather than phenomenological and to cast heterophenomenology as a study of fictions. Dennett proposes replacing phenomenology, or first person description of experience, with a conjoint first-third person approach: third person interpretation and causal explanation of first person description of experience *that may revise the first person description*.

Heterophenomenology is recommended as being 'neutral with regard to the debates about subjective versus objective approaches to phenomenology, and about the physical or nonphysical reality of phenomenological items' (95). The neutrality derives from the fact that Dennett conceives heterophenomenology as the study of a fictional narrative and of the fictional items and events posited therein. We do not typically debate the subjectivity or objectivity, or the physical as opposed to nonphysical reality of fictional items and events such as Sherlock Holmes' violin playing. Dennett recommends heterophenomenology, inso-

far as it is conceived as the study of a fictional narrative, on the ground that it circumvents analogous debates about conscious experience.

Moreover, on the ground that our conscious experiences are posits in fictional narratives, Dennett argues that it is possible that when we refer to our conscious experiences in the first person we are really referring to our brain events. When we talk about our experiences, brain events are the only real goings-on to which we could *really* be referring — that is, non-fictionally. Dennett supports this suggestion with an account of the interpretation of fiction that is controversial at best. His claim is that novels which can be understood in many respects as loosely veiled autobiographies are really *about* the loosely veiled rather than the fictional events. His justification for this claim is that since it is the real events 'that explain why *this* text got created' (365), it is *about* those events. Clearly, the literary analogy itself invites the charge of conflating causal explanation of the genesis of a fictional work with its interpretation. Nonetheless, this is Dennett's case for the *in principle possibility* that the referents of our claims about conscious experience are brain events.

He goes on to specify the facts that would justify us in claiming that this possibility is actually the case.

if we were to find real goings-on in people's brains that had *enough of the "defining"* properties of the items that populate their heterophenomenological worlds, we could reasonably propose that we had discovered what they were *really* talking about — even if they initially resisted the identifications. And if we discovered that the real goings-on bore *only a minor resemblance* to the heterophenomenological items, we could reasonably declare that people were just mistaken in the beliefs they expressed, in spite of their sincerity. (85, my added emphasis)

Dennett is proposing that (i) the *identification* of conscious (personal) experiences with (subpersonal) brain events is justified if brain events have 'enough of the "defining" properties' of conscious experiences; and that (ii) the elimination of conscious experiences as manifest at the first person viewpoint is justified if brain events have 'only a minor resemblance' to conscious experiences. So, for the sorts of destructive inferences Dennett wants to make, brain events must share 'enough of the "defining" properties' of conscious experiences while bearing only a minor resemblance to them. I doubt the coherence of this proposal because I doubt the coherence of the supposition that items which share some defining properties might bear only a minor resemblance to one another.

## **IV** The Plurality of Explanations

Be that as it may, what is clear is that this proposal departs from the view that in describing ourselves as experiencing persons we are making a distinctive kind of sense of ourselves. Or rather, it pays lip service to that view while denving its central commitments. Dennett's new framework does cast description of experience at the personal level as distinctive, but casting the distinctiveness as that of a *fictional* narrative repudiates the thesis that explanation at the personal level is necessary and ineliminable. We can let go of our fictions, however attached to them we may be and however important they may be to our practices. In contrast, personal level explanation has been argued to be ineliminable in the sense that without it our understanding would be incomplete, we would be unable to capture certain generalizations and facts. This point has often been made in terms of the notion of patterns, namely that there are certain patterns in our behavior and our interaction with the environment which can only be captured in concepts that pick out events at the personal, intentional, semantic level. Dennett has also advanced arguments of this kind.3

The first part of Dennett's theoretical framework — that dualism is just defeatism, so the brain must be the mind — doesn't even countenance the work, including Dennett's own, which argues that there is a necessary duality or plurality of distinctive explanations. (Indeed, by omitting any mention of this work, Dennett gives the nonprofessional reader, for whom the book is ostensibly written, the impression that duality of substance is the only issue.) Leaving this approach unmentioned misrepresents the current theoretical landscape. If this approach is correct, it would block the immediate inference from 'dualism is just defeatism' to 'the brain must be the mind.' The immediate inference is blocked by showing that a different consequence might follow from the denial of dualism. It may be argued that while dualism is defeatism, and the brain *does* provide the causal conditions for the mind, the brain *is not* the mind because to speak of and study minds — that is, to speak of and study ourselves as mindful or experiencing subjects — is to engage in a distinctive way of making ourselves and our behavior intelligible. Since

<sup>3</sup> See Dennett's 'Intentional Systems,' Brainstorms (Montgomery, VT: Bradford Books 1978); 'The Intentional Stance,' The Intentional Stance (Cambridge, MA: The MIT Press 1987); and especially 'Real Patterns,' The Journal of Philosophy 89 (1991). For an outstanding discussion of 'Real Patterns' and of considerations about patterns in theorizing about mind more generally, see John Haugeland's 'Pattern and Being,' in B. Dahlbom ed., Dennett and his Critics (Oxford: Basil Blackwell 1993).

facts figure in ways of making sense, insofar as there are distinctive ways of making sense, the facts articulated and explained therein cannot be mapped onto one another. Hence, insofar as facts about minds figure in a distinctive way of making sense of ourselves, they cannot be mapped onto facts about brains that figure in another way of making sense of ourselves

So, the last and most fundamental consideration in assessing Dennett's reductive framework is whether that framework needs to address the possible distinctiveness of the plurality of explanations of ourselves. Can a reductive framework be adequately supported without countering the case for distinctive explanations? Also, why might Dennett believe that he need not address those arguments?

These questions can be answered by reconstructing the implicit considerations behind Dennett's reductionism in Consciousness Explained. The reconstruction draws on the fact that Dennett's earlier work is shot through with tensions concerning the distinctiveness of explanations of ourselves. Hence it may be argued that the strands in Dennett's earlier work which tell against the distinctiveness of personal level explanations support the reductive framework of Consciousness Explained without needing to be rehearsed in that work. But are these earlier strands sufficient to obviate the need for addressing the distinctiveness of explanations in Consciousness Explained?

Dennett argues for the distinctiveness of personal level explanations in his work on (i) intentional systems and the intentional stance from which such systems come into view, (ii) semantic engines, (iii) persons and facts at the personal level. Interwoven with his development of these concepts are considerations which suggest that personal level explanations are different but not distinctive. The crucial strand is Dennett's conception of rationality: the rationality of persons, intentional systems, and semantic engines. His most explicitly developed argument is that the rationality of intentional systems is of the same kind as the rationality of design systems. This would mean that content attributions to intentional and design systems are governed by the same regulative principle rather than by distinctive ones. If these two types of explanations do not involve distinctive regulative principles, explaining something as an intentional system would not be distinctive from explaining it as a design system. Since Dennett explains consciousness by identifying conscious judgments with the states of the virtual machine, his account depends on the identifiability of intentional states with the states posited in a design explanation. Hence his project in Consciousness Explained would be underwritten insofar as his earlier work argues that the states posited in intentional and design explanations may be aligned and even possibly identified.

Let's begin with Dennett's considerations for the distinctiveness of intentional systems, semantic engines, and persons. Intentional systems are systems whose patterns of interaction with their environment are so complex that their behavior can only be predicted by attributing contentful or intentional states to them. Thus such intentional states are ascribed to the system as a whole. In contrast, the behavior of design and physical systems is explained in terms of their parts — functional subsystems and physical configurations — and the states of those parts. Hence, even though one and the same thing, metaphysically speaking, which is an intentional system is also a design system and a physical system, intentional states cannot be identified with functionally or physically individuated states. Dennett's insight that intentional states are states ascribed to the system as a whole in making sense of the system's patterns of behavior and interaction with its environment precludes the identifiability of intentional states with states of parts of the system (however those parts may be individuated).4

The point of Dennett's idea of semantic and syntactic engines is to distinguish two very different ways in which things may function. Semantic engines run on meanings, syntactic engines do not. Syntactic engines respond to physical or formal properties (simply, to non-semantic properties such as 'shapes' or connection strengths). Semantic engines respond to semantic properties. The point and promise of the idea of syntactic engines is that while such systems work purely non-semantically, their outputs and workings may be semantically interpretable nonetheless. We thus have a way of understanding how at least some physical things can produce semantically interpretable outputs. But this does not suggest that all semantic engines are to be explained reductively as identical to or constituted out of syntactic engines. The point of distinguishing semantic and syntactic engines is to allow for the possibility of identifying some physical systems as semantic engines and not only as syntactic ones. The distinction was originally conceived to distinguish the way in which we as conscious subjects respond to and deal with meanings while our central nervous systems do not. The need to appeal to the activities of our central nervous system in explaining the fact that we respond to meanings in no way fixes it that the explanation is to be reductive, reducing the semantic engine to the syntactic one. Moreover, the distinction in kind between semantic and syntactic engines remains even if one agrees with Dennett that all semantics, including that of semantic

<sup>4</sup> See John Haugeland's argument that elements of certain types of patterns are not identifiable with elements that can be individuated independently of their part in those patterns in 'Pattern and Being.'

engines, is a matter of interpretability. Interpretable semantic engines are such by virtue of responding to meanings while interpretable syntactic engines are such by virtue of responding to formal or physical features.

Dennett's distinction of facts at the personal level from those at the subpersonal level is intended to protect the events and features which figure in making sense of persons from identification with the events which figure in making sense of subpersonal parts — either physically or functionally individuated. Even though a physical whole is physically constituted out of its physical parts, some (physical) wholes — namely persons — are made intelligible in a way that is distinctive from the way in which their (physical) parts are made intelligible. Consequently, the states and events of the parts and of the whole referred to in those distinctive ways of making the parts and the whole intelligible do not stand in constitutive relations: the states of the parts do not constitute the states of the whole, just as the events involving the parts do not *constitute* the events involving the whole.

Dennett used to capture this point with the idea of switching topics: when we proceed from describing and explaining persons to explaining what transpires with their parts we switch subject matters.<sup>5</sup> If the subject matter of explanation at the personal level is distinct from the subject matter of subpersonal explanation, then discoveries at one level cannot force conceptual revision at the other level.

When we have said that a person's in pain, that she knows which bit of her hurts and that this is what's made her react in a certain way, we've said all that there is to say within the scope of the personal vocabulary. We can demand further explanation of how a person happens to withdraw her hand from the hot stove, but we cannot demand further explanations in terms of "mental processes." If we look for alternative modes of explanation, we must abandon the explanatory level of people and their sensations and activities and turn to the sub-personal level of brains and events in the nervous system. But when we abandon the personal level in a very real sense we abandon the subject matter of pains as well....

<sup>5</sup> It is interesting to note, and I grateful to William Seager for reminding me, that Donald Davidson has also made this point using the notion of switching the subject matter. See 'Mental Events' and 'Philosophy of Psychology,' reprinted in Essays on Actions and Events (Oxford: Clarendon Press 1980). In the latter essay Davidson writes: 'When we attribute a belief, a desire, a goal, an intention or a meaning to an agent, we necessarily operate within a system of concepts in part determined by the structure of beliefs and desires of the agent himself. Short of changing the subject, we cannot escape this feature of the psychological; but this feature has no counterpart in the world of physics' (230, my italics). This is especially noteworthy in light of the divergence between Dennett and Davidson that results from Dennett's understanding of rationality. I go on to examine this divergence in Davidson's and Dennett's positions.

Abandoning the personal level of explanation is just that: abandoning the pains and not bringing them along to identify with some physical event.<sup>6</sup>

Unlike the previous two sets of distinctions, the status of the distinction of personal from subpersonal facts is equivocal in the following sense. The question is whether the distinction of personal and subpersonal levels is self-standing or dependent on other more theoretical distinctions such as those of intentional and design systems or semantic and syntactic engines. But for our purposes it is not important whether we take Dennett to be articulating three sets of distinctions or two sets of distinctions with a third dependent set.

The concept of rationality is integral to each of these three sets of distinctions. The attribution of contentful states to persons, intentional systems, or semantic engines is governed by the regulative principle of rationality. Whether those attributions figure in distinctive explanations depends on how the regulative principle of rationality is conceived. We can understand the strain in Dennett's position by contrasting his conception of rationality with a stronger notion.

Intentional explanations are distinctive if the principle of rationality which regulates intentional-state attributions is conceived as a *constitutive ideal*. Donald Davidson's work exemplifies this strong position. According to Davidson, the regulative principle of rationality which governs attribution of intentional episodes is a *constitutive ideal* so that attributions are determined in terms of what ideally *ought to be* the case rather than by what generally *tends to be* the case. The idea is that in

<sup>6</sup> Content and Consciousness (London: Routledge & Kegan Paul 1969), 93-4; my added emphasis

<sup>7</sup> On one reading, Dennett's claim that examination of subpersonal facts switches the topic away from personal facts cannot stand on its own because it asserts from within the manifest image that manifest facts are part of a distinct subject matter and so are distinct from the facts revealed in scientific explanations. According to this interpretation, it is precisely because the assertion is made from within the manifest image that it cannot advance a claim about the relation of manifest and scientific facts. Hence the claim is taken to depend on how theoretical debates about the distinctness of intentional and design systems and semantic and syntactic engines turn out.

Alternatively, the distinctness of personal and subpersonal facts may be interpreted as making explicit something in the manifest image which is not dependent of how things turn out in the scientific image. Namely, manifest pain is a fact about the person as a whole. As such, the subject matter of pain is different from anything concerning parts of persons. On this reading, a claim about the manifest image made from *within* that image can have implications for the distinctness of manifest facts without depending on more theoretical debates.

forming and holding our beliefs we strive after an ideal and are held responsible to that ideal, so that a person's mental life does not have a structure where matters 'take their course.' At the first person viewpoint, what shapes my mental life is my attempt to live up to the dictates of rationality — so for example, what shapes the sorts of additions and revisions I make in my beliefs is the attempt to believe what I *ought*. From the third person viewpoint, when I attribute beliefs to someone else as others do to me. I do so on the basis of the understanding that I am making attributions to someone who is in some sense aware of his or her own responsibility for maintaining a rationally coherent mental life. In other words, we make attributions in the understanding that the structure of a person's beliefs, wishes, and suppositions has been informed by the attempt to believe and wish as one ought. In short, we make sense of persons as acting rationally and so the mental representations we attribute to them in doing so are rationally structured. To attribute mental representations is to construe the subject of those mental representations as a rational agent. These

In contrast, the attributions of contents which we might make in a subpersonal explanation are governed by a general tendency principle rather than a constitutive ideal. When our project is to explain the sub-personal causal conditions which make it possible that a person is the kind of being which can be understood as a rational agent with a subjective outlook, etc., we may individuate some of those conditions as content-bearing states and episodes. But the contents attributed in this project are of the kind that is attributable given the way things generally tend to happen as determined by causal principles and initial conditions. The way things generally tend to happen, as determined by causal principles and initial conditions, does not reflect what occurs in a structure 'in the attempt to live up to an ideal.' In short, the structure of contents that is attributed at the sub-personal level is not one for which the structural principle is the 'attempt to live up to an ideal.' This is why Davidson and others following him have claimed that a structure defined in terms of the constitutive principle of rationality has 'no echo in physical theory.'8

In contrast, Dennett holds that the rationality of intentional systems and stances is of a weaker (non-ideal) 'means-ends' variety. He rejects the idea that the regulative notion of rationality functions as a 'constitu-

are two faces of a single coin.

<sup>8</sup> Davidson, 'Psychology as Philosophy,' 231

<sup>9</sup> See 'Intentional Systems.'

tive ideal' on the ground that this would cast our rationality as perfect or as perfectly free and our mental lives as perfectly structured. It would cast us as perfect believers, always believing as we ought. Since we are clearly imperfectly rational, and the mental representations attributed in making sense of us do not hang together perfectly rationally, Dennett believes that it is incorrect to suppose that those attributions are governed by the normative *ideal* of rationality. Intentional episodes are attributed in an explanatory framework of concepts regulated by considerations of general tendency. According to Dennett, the intentional episodes attributed to intentional systems *approximate* episodes which stand in genuinely (in the sense of perfectly) rational relations.

I believe that Dennett's position is motivated by a misunderstanding of the idea that a regulative principle acts as a *constitutive ideal*, in particular, of the idea that the principle of rationality acts as a constitutive ideal in regulating mental content attributions. To make sense of something in terms of an ideal does not render what is thereby made intelligible ideal or perfect as well. For example, making sense of something as a logic proof attempt does not require making sense of it as valid. We can make sense of imperfect proof attempts as precisely that, but we do so by using an ideal structure which does not itself come in gradations or approximations (i.e., there is no such thing as a somewhat valid proof or an approximately valid proof in terms of which the order in an invalid proof attempt can be made intelligible).

Moreover, what is crucial is that the sort of order that is revealed in terms of a *constitutive ideal* answers to constraints that distinguish it from the sort of order revealed in terms of general tendencies. For example, 'discovery' of a certain heretofore unrecognized logical implication would require alterations in the contents to which we are committed at the first person and which we attribute at the third person viewpoint. Similarly, 'changes' in our understanding of what is rational or logical would require revisions. But these are not the sort of alterations required, or allowed, in a structure of states or episodes for which the structural principle is of a 'general tendency' type.

Hence I believe that Dennett's earlier work does not succeed in showing that intentional and design explanations are governed by the same type of regulative principle. Insofar as his work does not establish that the regulative principle governing all types of content attributions is the same, it does not establish that the states posited in those explanations

<sup>10</sup> See Elbow Room (Cambridge, MA: The MIT Press 1984), Ch. 2 'Making Reason

may be aligned or even identified. Thus Dennett's earlier work on the rationality of design and intentional systems does not underwrite his project in *Consciousness Explained*.

It is important to note that even if one agrees with Dennett that content-attributions to any finite system are regulated by a general tendency principle of rationality, that would not suffice to remove or bridge the other distinctions. It is not just the rationality informing semantics that makes responding to meanings distinctive from responding to formal properties. Hence, despite Dennett's gestures to the contrary, 11 the difference between semantic and syntactic engines is not removed by arguing that responses to meanings are only approximately rational. Similarly, considerations about rationality do not exhaust the distinction between personal and subpersonal facts, whether that distinction is considered to be self-standing or not. That pain is a fact about a person as a whole rather than a fact about her parts is not affected by the strength of one's notion of rationality. Thus, at most, one might understand Dennett's earlier work as poised between the distinctiveness of syntactic and semantic, and personal and subpersonal explanations on the one hand and the potential alignment of intentional and design explanations on the other. As such, the earlier work suffices neither to ground the reductive framework of Consciousness Explained nor to obviate the need to discuss these issues.

#### V The Destructive Inferences

Dennett needs the reductive framework in order to justify using his account of the subpersonal virtual machine to draw inferences about conscious experience. According to that account, conscious experience is neither qualitatively nor temporally continuous and determinate. That is, our conscious experience does *not* consist of a continuous, fully determinate stream of qualitative states or episodes. This is because we find neither temporal continuity, temporal determinacy nor qualitative continuity when we look to the relevant subpersonal processes and states.

The states which make up the stream of our conscious experience do not have a *qualitatively continuous* character because the states of the virtual machine with which they are identical are discrete in function and realization. Let's follow Dennett and consider perceptual episodes as our paradigm case of the sort of conscious experiences which have a

<sup>11</sup> See Elbow Room, Ch. 2 especially 28-9.

qualitative nature. According to Dennett, subpersonal perceptual processes have the function of making discriminations. The contents of these processes (or states) are what they discriminate. For example, processes which discriminate a red external object thereby have or fix the content red. The functional role of discriminatory states does not require them to have a qualitative aspect in order to carry out their function. Hence there is no need or warrant for attributing such an aspect to them. Just as the discriminatory states of a computer do not need to have qualitative aspects in order to make color discriminations, so the discriminatory states of an experiencing human being need not have a qualitative aspect in fulfilling their function (374). Hence, though perceptual experience seems to be one way or another — there is something it seems like to be perceiving a blooming lilac or a cardinal in song — Dennett holds that there is no causal basis for the seeming and so there is no seeming, not really. More generally, even though conscious experiences may seem to us to be one way rather than another or to have a qualitative aspect, there is no seeming, not really. Dennett is denying that there is any causal basis for the seeming, for the qualitative nature of much conscious experience, and since there is no causal basis there is no seeming, not really.

... the seeming isn't rendered at all....

There is no such phenomenon as really seeming — over and above the phenomenon of judging in one way or another that something is the case.

There seems to be phenomenology. That's a fact that the heterophenomenologist enthusiastically concedes. But it does not follow from this undeniable, universally attested fact that there really is phenomenology. (355, 364, 366)

Dennett's metaphoric way of denying that there is anything it seems like to be conscious is that conscious experience is really all tell and no show. The metaphor likens the content of conscious experience to linguistic content while denving that it might be akin to pictorial content. At least apparently, pictures bear contents that are continuous in character or nature by means of the continuous properties of their vehicles. Dennett directs us to think of something like the continuous expanses of pigment which would be used to realize a pictorial content such as: the blooming lilac. In contrast, at least apparently, linguistic contents are not continuous in character and they do not need to be borne by continuous properties of their vehicles. Here we are to think of something like the concatenation of discrete symbols which are used to bear a linguistic content such as: 'the blooming lilac.' Dennett's point is that experiential content isn't really continuous in nature, like imagistic content, because the causal basis of conscious experience involves discrete vehicles namely discrete discriminatory states. He is assuming that since the causal basis involves discrete states or vehicles, the content attributable to those vehicles must be of an essentially 'discrete' kind as well. The

implication of denying that conscious experience might show while telling is that there is no real difference between imaging blooming lilacs and thinking *that* lilacs are blooming, between hearing a cardinal's song, and thinking *that* a cardinal is singing.

Just as our conscious experiences seem to have a continuous qualitative character, they also seem to make up a continuous stream across time. However, there is no causal basis for the apparent temporal continuity since the causal basis of conscious experience consists of discrete states. The causal basis of consciousness also belies its seeming temporal determinacy. Conscious experience seems to be determinate at any given moment. What my conscious experiences are from moment to moment, their nature (for example whether they are perceptions or memories) and their order all seem determinate. There seem to be facts of the matter about what I experience from moment to moment — whether I am perceiving blooming lilacs or recalling a cardinal's song — and about the order of those experiences. However, according to Dennett there are no such facts of the matter about conscious experience because conscious states are states of the serial virtual machine, and the latter are subpersonal states whose selection depends on their functional effects on behavior, effects that can be changed or canceled out by other subsequent states.

There is no reality of conscious experience independent of the effects of various vehicles of content on subsequent action (and hence, of course, on memory). (132)

The denial of determinate facts of conscious experience follows from Dennett's account of the selection process by which multiple parallel states are selected into a serial virtual machine. According to the Pandemonium model, multitudes of discriminations or content-fixations are continuously interacting, compounding and/or fading in parallel. There is nothing to distinguish among these multitudinous parallel contentfixations aside from their subsequent effects on other competitors in the Pandemonium. Content-fixations form into more complex content-coalitions depending on their effects. A coalition of contents wins out in the Pandemonium and is a state of the virtual serial machine insofar as it affects behavior. To affect behavior and win, a coalition of contents needs to 'answer' a 'question' or probe. Winning coalitions answer one's own questions, those from others or from the environment. Consequently, the selection of content is determined by the nature and timing of the probes. Since the selection process is one of continuous probing from a variety of sources into a multitudinous pool of potential answers, varying discrepant answers may be selected across time. This process allows for the alteration of the effects of one winner by a subsequent one. Hence there cannot be a fact of the matter about conscious experience at a time, since a subsequent winning content may change precisely that.

To see how discrepant answers may overwrite their predecessors completely (even overwriting the very occurrence of their predecessors) one must keep in mind the content/vehicle distinction as it applies to subpersonal discriminatory states. A functionally specifiable sequence of states which makes up the virtual (serial) machine is a sequence of physical states or 'vehicles' to which certain contents are attributable. Simply put, the contentful states of the virtual machine consist of physical vehicles which bear certain contents by virtue of the fact that there are grounds for attributing those contents to them. Consequently, when it comes to the timing of such contentful physical states, the time of the representing and the represented time may be distinct. That is, the time of the representing is the time at which the representing vehicle occurs, and it may be distinct from the represented time which is the time specified in the content attributed to and carried by that vehicle. Simply, the represented time of a certain content may differ from the time at which the vehicle bearing that content occurs. Since it is the *contents* of the states of the virtual machine that are identical with the judgments making up our stream of consciousness, the time at which we experience a judgment is the represented time, the time specified as part of the content of that judgment (rather than the time at which the vehicle occurs). Thus, a state of the serial machine may erase and overwrite the effects of an earlier state by representing its own time as the time of the earlier state (that is, as the time that the earlier state represented as its own).

Finally, according to the Pandemonium model, processes that contribute to a content-fixation and those that modify it immediately afterwards do not have distinguishing effects. At the personal level of conscious experience it seems that such alterations can be distinguished into the vagaries of memory on initial, perhaps perceptual, experience. However, at the micro time scale of our subpersonal parallel processes, there is only continuous interaction among competing content-fixations. But if the processes that construct perceptions and those that reconstruct them, perhaps revising them, are not distinguishable functionally at the micro level, there is no causal basis for distinguishing them at the macro level. Hence, there is no causal basis (independent of the contingencies of probing necessary for behavioral effects) for a principled difference between the initial construction of experience — or perception — and its reconstruction — or memory.

... this is the fundamental implication of the MD [Multiple Drafts] model — if one wants to settle on some moment of processing in the brain as the moment of consciousness, this has to be arbitrary. One can always "draw a line" in the stream of processing in the brain, but there are no functional differences that could motivate declaring all prior stages and revisions to be unconscious or preconscious adjustments, and subsequent emendations to the content (as revealed by recollection) to

be post-experiential memory contamination. This distinction lapses in close quarters. (125)

In sum, Dennett argues that the Multiple Drafts Model shows that there is no fact of the matter about the content of conscious experience — since conscious experiences are identical to the contentful states of the virtual serial machine which may alter and even overwrite each other's effects, and amongst which no principled differences can be drawn. Hence, even though there are always facts of the matter about the vehicles (physical configurations) which bear the contents of the serial virtual machine, there are nonetheless no fixed, determinate facts about the contents.

However, these destructive inferences do not follow *just* from placing the Multiple Drafts Pandemonium Model in Dennett's reductive framework. Certain other substantive assumptions are also required. Dennett neither identifies nor supports these additional substantive theses. I will highlight these further theses in order to show their lack of warrant.

To conclude that conscious experience is all tell and no show because subpersonal contents tell rather than show, one needs to hold that personal and subpersonal contents must be of the same type. That is, one would only make this inference if one believed that subpersonal contents explain personal, conscious contents insofar as they are of the same type and that subpersonal contents dictate which type that is since they are the contents attributed to the physically real causal basis (namely, our brain processes). To conclude that there is no fact of the matter about the nature or temporal order of conscious experience because the distinction between initial content-fixation and subsequent alteration lapses at the micro time intervals of subpersonal functioning, one needs to hold that the same features must obtain at all time scales or all explanatory levels.

I suggest that just as we need not hold that physical objects are spatially 'gappy' rather than solid because of the 'gappiness' between atoms at the micro level; so we need not hold that conscious experience is 'gappy' rather than temporally and qualitatively continuous because of the 'gaps' between subpersonal states. But at the very least, we need to recognize the issue.<sup>12</sup>

<sup>12</sup> The thesis that manifest features of conscious experience are not real unless there are corresponding features at the micro time durations of the causal basis is analogous to the thesis that manifest features of physical objects are not real because they are lacking at the micro spatial scale. The latter thesis cannot be claimed to be the received view about spatial properties in the philosophy of science. Rather, it is countered by a 'plurality of distinctive explanations' type of approach. I mention the analogy so as to point to a lacuna in Dennett's argumentation. Dennett's

So what need subpersonal conditions be like in order to explain the enabling conditions of experiential content? Any attempt to provide a causal explanation of the subpersonal conditions that enable experiential content faces this issue. How 'close' need the similarity be between experiential content and the content attributed to the virtual machine in order for the latter to fulfill its explanatory role? Let's consider one obvious approach, namely that subpersonal processes need to bear contents of the same type as the personal contents being explained. I have indicated that Dennett's 'all tell and no show' case against anything like 'qualitative' content at the personal level makes this very assumption. But the Multiple Drafts Model undermines the motivation or need for that assumption. A Multiple Drafts or Parallel Distributed Model of subpersonal functioning shows that there need be no obvious match at all between subpersonal processes and events and those of the experiencing person. This is also precisely what commitment to the plurality of distinctive explanations entails. If personal and subpersonal explanatory projects are distinctive, there is no theoretical reason to expect that contentful episodes at the personal level need to be 'matched' at the subpersonal level by episodes and types of contents borne thereby. For example, if one holds that subpersonal conditions stand in a causally enabling relation to the attributions that are made to the person, then the nature of those subpersonal conditions need not stand in any obvious correspondence or similarity to personal level content. Theoretical commitment to the distinctiveness of personal and subpersonal explanatory projects directs one to be open concerning the kind of content, if any, which will figure in the best explanation of the subpersonal enabling conditions of personal content.

Indeed, if it turns out that the events, states and processes (and types of contents) that figure in our best subpersonal explanation are different in type from the personal episodes whose enabling conditions are being thus explained, this would not render the personal episodes fictional. For example, while it might be appropriate to make sense of a subject as *noticing* something, the subpersonal explanation need not posit a corre-

discussion focuses on the verificationist (or operationist) assumption that a feature is not real unless it makes a discernible difference in behavior. However, that assumption does not come into play unless one has already accepted the thesis that features and distinctions must persist at micro temporal durations. Dennett does not raise this more fundamental issue, thus giving the reader the impression that there is only one rather than two controversial matters here. If one does not accept the thesis that features must persist across spatial and temporal scales, there is no need to get exercised about Dennett's verificationism.

sponding identifiable event or episode of noticing with a corresponding type of content.<sup>13</sup>

In examining Dennett's theoretical framework, I discussed his pre-Consciousness Explained insight that content attributions to the experiencing subject and to the virtual machine could not be aligned or identified because the former are attributions to the whole while the latter are attributions to the parts. This insight is augmented in Consciousness Explained with Dennett's claim that the multiple subpersonal discriminations or content-fixations need not be expressible in 'propositional form.' In other words, content attributed in a subpersonal explanation might be utterly unlike the propositionally expressible contents that figure in making ourselves intelligible as experiencing persons. But if this is so, then why would Dennett also try to maintain that subpersonal discriminations are to be likened to telling (i.e., to linguistic representation)?

If the best subpersonal explanation is utterly unlike phenomenological description — which is the overarching conclusion Dennett wants to establish — and if what is going on are parallel processes implementing a serial virtual machine — which is Dennett's favored empirical model - why do we need to think of such processes as telling rather than showing? Subpersonal content-fixations might involve a third different type of content which nonetheless provides the subpersonal conditions for the sorts of contents we experience. Moreover, if we need not think of the subpersonal events as telling rather than showing (since we need not think of them as even propositionally expressible), there is no reason to think that the personal events really tell rather than show as well. In short. Dennett's case against qualia loses sight of his own insight that subpersonal content may be so unlike personal content as to be inexpressible in propositional form. In the previous section I examined the broader theoretical context of losing the difference between subpersonal and personal content from view and the fact that its payoff is reductionism. At this point we are in a position to note that Dennett's loss of sight prevents us from making the most of the new explanatory possibilities made available in his Multiple Drafts/Pandemonium Model. More generally, we can appreciate that hankering for any 'tight' fit between our personal and subpersonal explanations blocks the explanatory possibilities opened by Parallel Distributed Processing Models.

<sup>13</sup> This example is John McDowell's; see his 'The Concept of Perceptual Experience,' Philosophical Quarterly 44 (1994).

#### Conclusion

'The Message is: There is no *Medium*.' So has Dennett summed up his objective in *Consciousness Explained*. Whether this is a fair summation depends on how one takes the notion of 'medium.' To conclude, let me discuss how my examination of the reductionism of this work shows that, in one crucial sense of 'medium,' *Consciousness Explained* delivers precisely the opposite message.

We have examined the thesis that conscious experiences are not realized in a medium in the sense of something like pigment (or in more technical language, that they are not realized or borne by continuous properties of their vehicles) and that hence they are not qualitative in nature. Dennett is also concerned to deny the me in medium, a thread I have neglected for lack of space. To deny the me in medium is to deny that there is a causal basis for the seemingly unified, directed, intentional nature of conscious experience — that is, of at least most of what we consciously think, say and do. Dennett denies that there is a causal basis for the apparently directed nature of our words and actions by sketching the undirected nature of the process which selects the contentful states of the virtual machine. Since the states of the virtual machine are selected in an undirected process from a Pandemonium of competitors, our thought, talk and action do not originate in a directed way. In other words. Dennett holds that if there is no functional centre to the brain, there is no unified, directed origin for our thought, talk and action. There is no central meaner, central intender or centre for any other function since there is only undirected selection from Pandemonium. (At most, there is a centre of 'narrative gravity.')

It may be correct that consciousness is not realized by a medium or pigment. (I have suggested that subpersonal conditions need not be explained as 'pigment' even while we continue to countenance the qualitative nature of conscious states.) And there are interesting senses in which there is no central meaner or source of intentions, not if they involve 'pre-linguistic' intentions. Insofar as these conceptions ought to be discarded, it is important to have an empirical model that shows that we can do without them. (That is, it is important to have an empirical model that shows that we can explain the subpersonal conditions which enable conscious experience without positing 'pigment' or processes which would provide the causal basis for pre-linguistic intentions.) The Pandemonium Multiple Drafts Model is one such model, and perhaps

<sup>14</sup> Daniel Dennett, 'The Message is: There is no Medium,' Philosophy and Phenomenological Research 53 (1993)

in its refined version it will be the best model of the subpersonal causal conditions of conscious experience.

Yet there is another sense in which conscious experience might have no medium that Dennett forsakes. Though I cannot develop this point in any length here. I believe that conscious experiences are contentful episodes which do not have vehicles. 15 Making ourselves intelligible as experiencing persons involves ascribing contents to ourselves — that is, ascribing only contents without vehicles or vehicle-less contents. For example, when I make sense of someone's actions by attributing to her the conscious or occurrent belief that 'It will probably rain,' making sense of her in this way does not involve attributing a vehicle which bears that content. The conscious contents we attribute in making sense of people are vehicle-less, and so conscious experiences do not have a medium in the sense that they are not carried by vehicles. On this view, one of the more interesting distinguishing features of the two explanatory projects we need to pursue in order to understand ourselves — the projects of making ourselves intelligible as experiencing persons and of explaining the subpersonal causal conditions that make it possible that we can be made sense of as persons — is that the former posits vehicle-less contents while the latter posits contents borne by vehicles.

By identifying contents at the personal level with subpersonal contents (the contents of the virtual machine which do have vehicles), a reductive framework like Dennett's provides contents at the personal level with vehicles and so, in this sense, with a medium. I believe that this constitutes the continuing appeal of reductionist explanations. Representations with which we are familiar, such as images and spoken or written linguistic expressions, all involve vehicles. Hence it may seem that for there to be contents there must be vehicles, since it is vehicles which bear contents. And so the idea of vehicle-less contents may seem mysterious or unstable and requiring further elucidation — elucidation which would show that seemingly vehicle-less contents are borne by vehicles after all.

The theoretical drive to explain all contents as carried by vehicles lies on a direct collision course with the suggestion that we can make sense of some very complex chattering beings by attributing experiential (or vehicle-less) contents to them. One way to make sense of the Dennettian trajectory through the space of explanations is in terms of this collision course.

<sup>15</sup> For a detailed presentation of this thesis, see my 'The Vehicle-less nature of experiential content.'

In urging that intentional contents are attributed to an Intentional System as a whole, pre-Consciousness Explained Dennett was suggesting that to make something intelligible as an Intentional System is to attribute to it contents without vehicles. In Consciousness Explained, Dennett consistently neglects any mention of a possible distinction between the vehicles that bear contents at the subpersonal level — states of content-fixation — and personal conscious content. Without addressing the issue explicitly, Consciousness Explained simply treats the vehicles of subpersonal content — namely states of content-fixation — as the vehicles of our conscious personal content.

A 'structural' perspective on this trajectory provides some more insight. Dennett's explanatory strategy has consistently been to 'develop an account of content that is *independent of* and *more fundamental than* consciousness ... and second, to build an account of consciousness on that foundation' (457). This strategy runs into trouble because of differences in the first and third person perspectives on content and consciousness. Dennett attempts to resolve that trouble by moving from a nonreductive account of intentional content to a reductive account of conscious (including contentful) states. He does so by providing an account of the subpersonal vehicles of conscious contents. But, a nonreductive account of intentional content and a reductive account of conscious content do not mix.

I introduced this tension in Dennett's approach to conscious experience with the observation that the manifest image of ourselves doesn't simply posit conscious (including contentful) experience as a way of making sense of one another but is itself consciously experienced at the first-person perspective. We *consciously experience* ourselves as consciously experiencing. The problem, as we noted, is to understand how this phenomenon might be akin to an explanatory posit. Let me elaborate.

From the third person theoretical standpoint we can conceive that experiencing ourselves as *thinking* experiencing persons could be a way of understanding into which we are acculturated. We can conceive that contentful mental episodes are analogues of explanatory posits in terms of which we not only grasp the mental lives of others but which articulate our own.<sup>17</sup> So it seems that this approach to the contentful or intentional nature of our mental lives can work at both the third and first person

<sup>16</sup> I present a detailed argument for this claim in 'The Vehicle-less Nature of Experiential Content.'

<sup>17</sup> See Wilfrid Sellars, 'Empiricism and Philosophy of Mind.'

perspectives. Similarly, at the third person viewpoint we can conceive that experiencing ourselves as thinking experiencing persons is a way of understanding into which we are acculturated. Perhaps the narratives which we learn to tell about ourselves to explain our behavior involve the idea that our contentful mental episodes are experienced. But at the first person perspective it is not quite so clear how to conceive of the experiential or conscious nature of our mental episodes as a feature of episodes we ascribe to ourselves. How could their conscious nature be something like a theoretical posit which we apply not only to others but to ourselves as well? The conscious or experiential aspect of our mental lives seems to be part of their very nature and so part of our nature.

Given that we are biological organisms made of physical processes governed by natural laws, it might seem more fitting to explain the conscious, experiential aspect of our mental lives in terms of those processes. Since our conscious episodes figure in the narratives we spin about ourselves, the temptation arises to align these explanations onto one another, matching each conscious episode and its features to some sequence or organization of physical processes. Since it is the same episodes which are conscious and which figure in our narratives, it may seem that there is a precise alignment of explanations to be had. If one takes this course, then since it is the physical processes which are metaphysically real, it will seem that any lack of 'fit' discovered between the nature of the physical processes and the features of the episodes posited in our narratives is to be resolved in favor of the metaphysically real processes.

This is, roughly, the trajectory from a Dennettian account of intentional content to the Dennettian account of conscious experience. Having contributed subtle and sophisticated conceptual tools for understanding ourselves as thinking persons and for understanding ourselves in terms of a plurality of distinctive explanations, Dennett turns (or rather returns) to the conscious nature of our thinking lives. And in so doing, he loses sight of the implications and requirements imposed by the very tools he has fashioned. That is, in attempting to explain the first person experiential nature of our mental lives, Dennett forsakes many of the theoretical commitments of his third person approach to the contentfulness of our mental lives.

The interesting twist, which makes Dennett's reductive alignment 'strained,' is his attempt to use the distinctiveness of making sense of ourselves as conscious to eliminate aspects of what it is to be conscious. Dennett holds that our concepts and the associated social or cultural norms play an integral role in the fact that there is consciousness and in what it is like. But he does so in order to eliminate precisely those aspects that are a function of concepts. In this way, Dennett's position appears to accommodate a pivotal commitment of the distinctive explanations

approach: it appears to countenance the norm-governed and hence social nature of conscious experience. But of course in holding that the aspects of consciousness which depend on our conception of it are eliminable, he denies the substance of the thesis that consciousness is, at least in part, an ineliminably norm-governed, social phenomenon. In this way, Dennett's account of consciousness appears to remain consistent with his earlier work on intentional content while departing in substance.

On the view of consciousness I will develop in this book, it turns out that consciousness, like love or money, is a phenomenon that does indeed depend to a surprising extent on its associated concepts. Although, like love, it has an elaborate biological base, like money, some of its most significant features are borne along on the culture, not simply inherent somehow, in the physical structure of its instances.... If everyone forgot what money was, there wouldn't be any money anymore; there would be stacks of engraved paper slips, embossed metal disks, computerized records of account balances, granite and marble bank buildings — but no money: no inflation or deflation or exchange rates or interest — or *monetary value*. The very property of those variously engraved slips of paper that explains — as nothing else could — their trajectories from hand to hand in the wake of various deeds and exchanges would evaporate.... So if I am right, and if I succeed in overthrowing some of those concepts, I will threaten with extinction whatever phenomena of consciousness depend on them. (24, my change of order of the text)

But what, if anything, might require us to discard aspects of our conception of consciousness? Might physical or biological or computational discoveries require us to do so? While consciousness, like money, might disappear if the associated concepts and social practices on which it depends were to disappear, the elimination of those concepts and practices would not be driven by discoveries about the nature of the physical basis (or realization). In the case of both love and money, the requisite concepts and practices are not beholden exclusively or primarily to facts about the physical basis. Let's consider the phenomena of love, since consciousness is more like love in having 'an elaborate biological base.'

The disappearance of chival rous love and its impossibility in our culture are not due to facts about its biological basis. Whatever the biological basis of love may be, we can be sure it has remained fairly constant across the seven hundred years separating us from Medieval courtly lovers. Cultural factors drive the changes in biologically based, culturally shaped phenomena. To see that this is true, we may try out Dennett's alternative view — that discoveries about the biological basis warrant or require changes in culturally shaped, biologically based phenomena — in the case of love. Mightitbe a discovery of brain science that one cannot really love the wrong person or that there is no fact of the matter about the timing of a love? Might one argue that love cannot be blind or induce rose-colored perception on

the grounds of discoveries about its biological base? Yes, if we reflect Dennett's work on consciousness back onto the case of love, these are *precisely* the sort of revisions to which his approach is committed. I suggest that we are justified in rejecting revisions so grounded, in the case of consciousness just as in the case of love.

Received: October, 1994

SONIA SEDIVY University of Toronto Toronto, ON Canada M5S 1A1