

and plausibility of very complex interactive information processes playing a role in the production of cognitive events so swift as to be atomic to introspection.

At last I turn to the important question. Suppose that AI is viewed as I recommend, as a most abstract inquiry into the possibility of intelligence or knowledge. Has it solved any very general problems or discovered any very important constraints or principles? I think the answer is a qualified yes. In particular, I think AI has broken the back of an argument that has bedeviled philosophers and psychologists for over two hundred years. Here is a skeletal version of it: *First*, the only psychology that could possibly succeed in explaining the complexities of human activity must posit internal representations. This premise has been deemed obvious by just about everyone except the radical behaviorists (both in psychology and philosophy—both Watson and Skinner, and Ryle and Malcolm). Descartes doubted almost everything

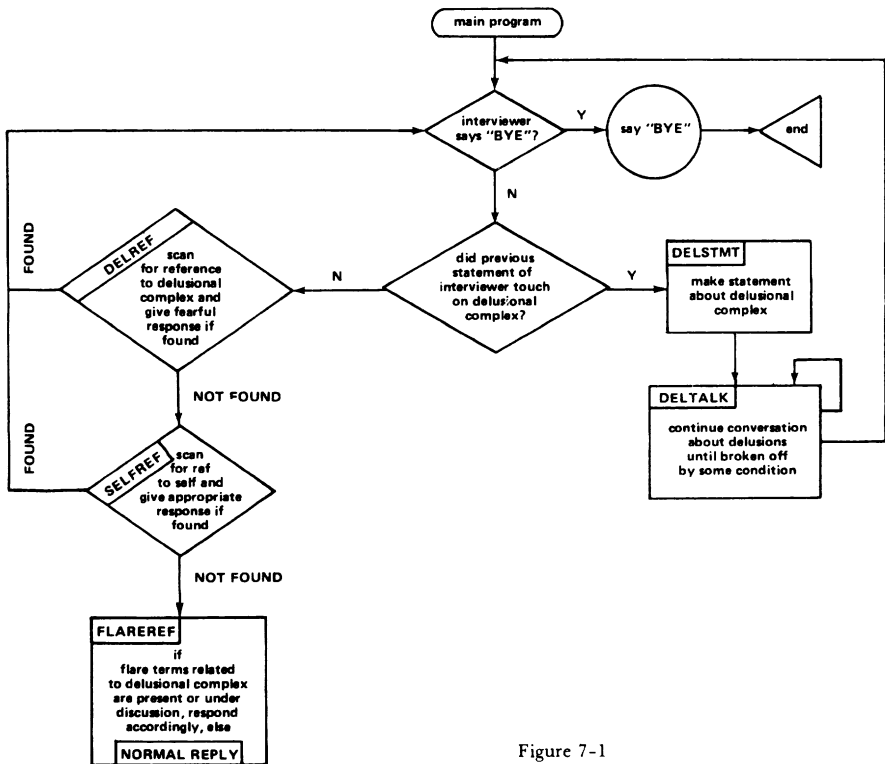


Figure 7-1

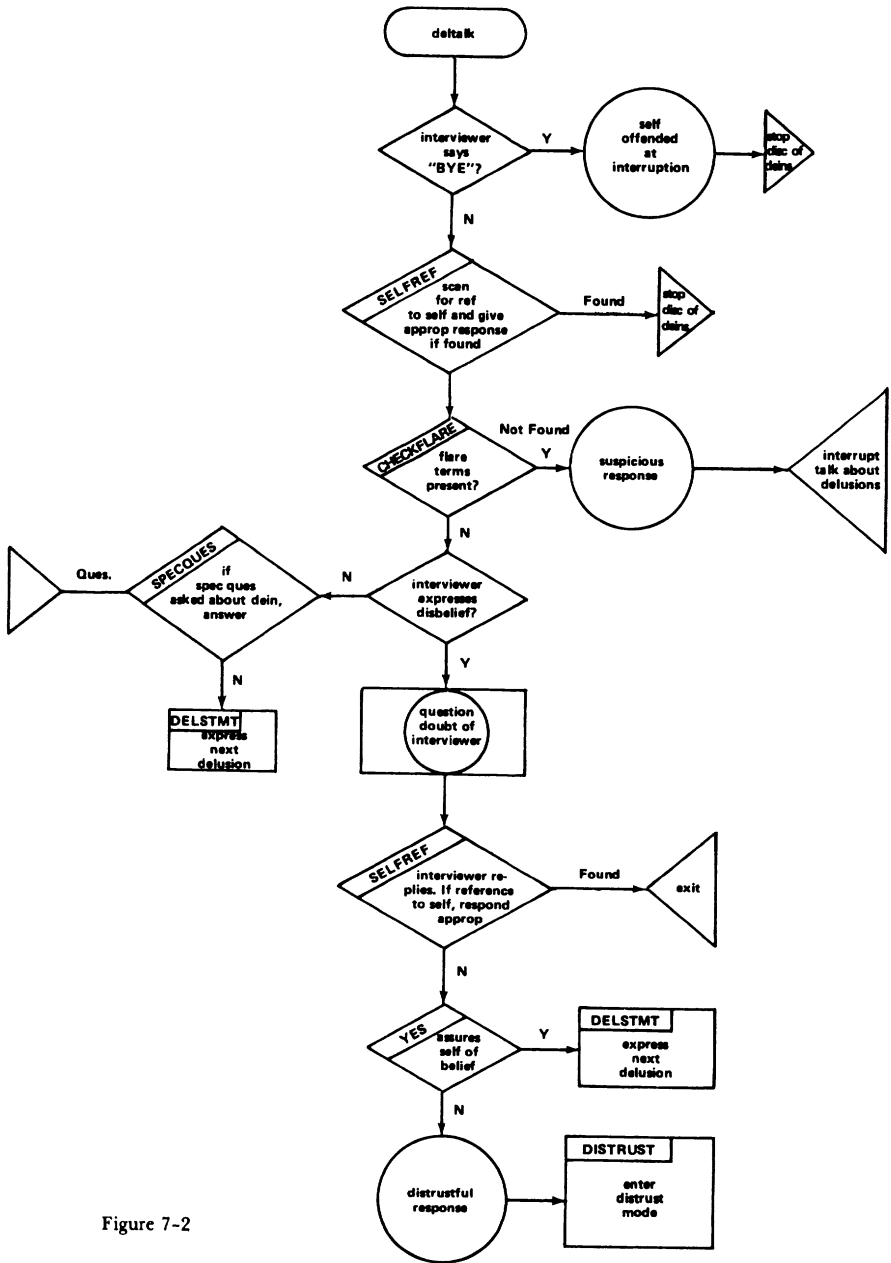


Figure 7-2

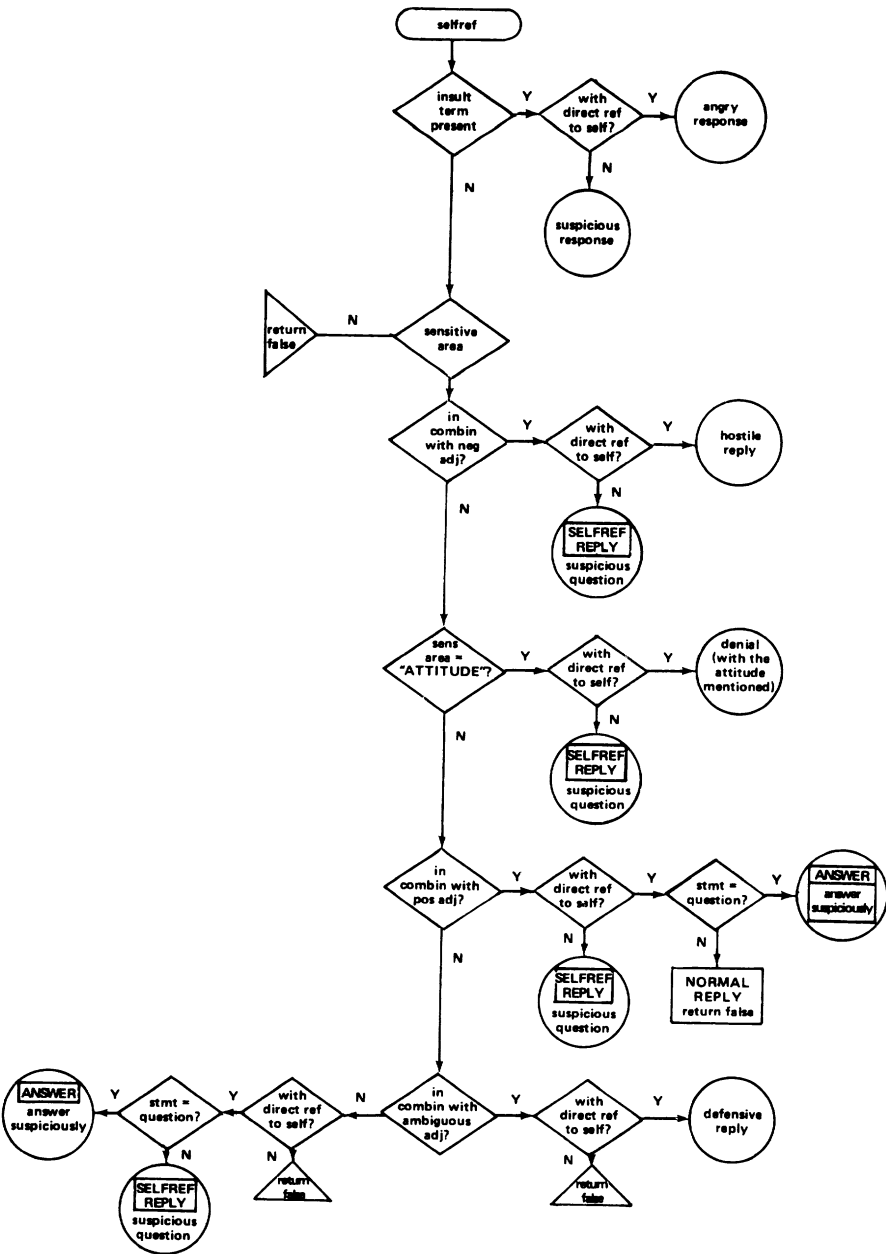


Figure 7-3

*but* this. For the British Empiricists, the internal representations were called ideas, sensations, impressions; more recently psychologists have talked of hypotheses, maps, schemas, images, propositions, engrams, neural signals, even holograms and whole innate theories. So the first premise is quite invulnerable, or at any rate it has an impressive mandate (see Chapter 6). But, *second*, nothing is intrinsically a representation of anything; something is a representation only *for* or *to* someone; any representation or system of representations thus requires at least one *user* or *interpreter* of the representation who is external to it. Any such interpreter must have a variety of psychological or intentional traits (see Chapter 1): it must be capable of a variety of *comprehension*, and must have beliefs and goals (so it can *use* the representation to *inform* itself and thus assist it in achieving its goals). Such an interpreter is then a sort of homunculus.

Therefore, psychology *without* homunculi is impossible. But psychology *with* homunculi is doomed to circularity or infinite regress, so psychology is impossible.

The argument given is a relatively abstract version of a familiar group of problems. For instance, it seems (to many) that we cannot account for perception unless we suppose it provides us with an internal image (or model or map) of the external world, and yet what good would that image do us unless we have an inner eye to perceive it, and how are we to explain *its* capacity for perception? It also seems (to many) that understanding a heard sentence must be somehow *translating* it into some internal message, but how will this message in turn be understood: by translating it into something else? The problem is an old one, and let's call it *Hume's Problem*, for while he did not state it explicitly, he appreciated its force and strove mightily to escape its clutches. Hume's internal representations were impressions and ideas, and he wisely shunned the notion of an inner *self* that would intelligently *manipulate* these items, but this left him with the necessity of getting the ideas and impressions to "think for themselves". The result was his theory of the self as a "bundle" of (nothing but) impressions and ideas. He attempted to set these impressions and ideas into dynamic interaction by positing various associationistic links, so that each succeeding idea in the stream of consciousness dragged its successor onto the stage according to one or another principle, all without benefit of intelligent *supervision*. It didn't work, of course. It couldn't conceivably work, and Hume's failure is plausibly viewed as the harbinger of doom for any remotely analogous enterprise. On the one hand, how could *any* theory of psychology make sense of representations that *understand themselves*, and on the other, how could *any*

theory of psychology avoid regress or circularity if it posits at least one representation-understander in addition to the representations?

Now no doubt some philosophers and psychologists who have appealed to internal representations over the years have believed in their hearts that somehow the force of this argument could be blunted, that Hume's problem could be solved, but I am sure no one had the slightest idea *how to do this* until AI and the notion of data-structures came along. Data-structures may or may not be biologically or psychologically realistic representations, but they are, if not living, breathing examples, at least clanking, functioning examples of representations that can be said in the requisite sense to understand themselves.\*

How this is accomplished can be metaphorically described (and any talk about internal representations is bound to have a large element of metaphor in it) by elaborating our description (see Chapter 5) of AI as a top-down theoretical inquiry. One starts, in AI, with a specification of a whole person or cognitive organism—what I call, more neutrally, an intentional system (see Chapter 1)—or some artificial segment of that person's abilities (e.g., chess-playing, answering questions about baseball) and then breaks that largest intentional system into an organization of subsystems, each of which could itself be viewed as an intentional system (with its own specialized beliefs and desires) and hence as formally a homunculus. In fact, homunculus talk is ubiquitous in AI, and almost always illuminating. AI homunculi talk to each other, wrest control from each other, volunteer, sub-contract, supervise, and even kill. There seems no better way of describing what is going on.<sup>11</sup> Homunculi are *bogeymen* only if they duplicate *entire* the talents they are rung in to explain (a special case of danger (1)). If one can get a team or committee of *relatively* ignorant, narrow-minded, blind homunculi to produce the intelligent behavior of the whole, this is progress. A flow chart is typically the organizational chart of a committee of homunculi (investigators, librarians, accountants, executives); each box specifies a homunculus by prescribing a function *without*

\*Joseph Weizenbaum has pointed out to me that Turing saw from the very beginning that computers could in principle break the threatened regress of Hume's Problem, and George Smith has drawn my attention to similar early wisdom in Von Neumann. It has taken a generation of development for their profound insights to be confirmed, after a fashion, by detailed models. It is one thing—far from negligible—to proclaim a possibility in principle, and another to reveal how the possibility might be made actual in detail. Before the relatively recent inventions of AI, the belief that Hume's Problem could be dissolved somehow by the conceptual advances of computer science provided encouragement but scant guidance to psychologists and philosophers.

*saying how it is to be accomplished* (one says, in effect: put a little man in there to do the job). If we then look closer at the individual boxes we see that the function of each is accomplished by subdividing it via another flow chart into still smaller, more stupid homunculi. Eventually this nesting of boxes within boxes lands you with homunculi so stupid (all they have to do is remember whether to say yes or no when asked) that they can be, as one says, "replaced by a machine". One *discharges* fancy homunculi from one's scheme by organizing armies of such idiots to do the work.

When homunculi at a level interact, they do so by sending *messages*, and each homunculus has representations that it uses to execute its functions. Thus typical AI discussions *do* draw a distinction between representation and representation-user<sup>12</sup>: they take the *first step* of the threatened infinite regress, but as many writers in AI have observed,<sup>13</sup> it has gradually emerged from the tinkering of AI that there is a trade-off between sophistication in the representation and sophistication in the user. The more raw and uninterpreted the representation—e.g., the mosaic of retinal stimulation at an instant—the more sophisticated the interpreter or user of the representation. The more interpreted a representation—the more *procedural* information is *embodied in it*, for instance—the less fancy the interpreter need be. It is this fact that permits one to get away with *lesser* homunculi at high levels, by getting their earlier or lower brethren to do some of the work. One never quite gets *completely* self-understanding representations (unless one stands back and views all representation in the system from a global vantage point), but all homunculi are ultimately discharged. One gets the advantage of the trade-off only by sacrificing versatility and universality in one's subsystems and their representations,<sup>14</sup> so one's homunculi cannot be too versatile nor can the messages they send and receive have the full flavor of normal human linguistic interaction. We have seen an example of how homuncular communications may fall short in SHRDLU's remark, "Because you asked me to." The context of production and the function of the utterance makes clear that this is a sophisticated communication and the product of a sophisticated representation, but it is not a full-fledged Gricean speech act. If it were, it would require too fancy a homunculus to use it.

There are two ways a philosopher might view AI data structures. One could grant that they are indeed self-understanding representations or one could cite the various disanalogies between them and prototypical or *real* representations (human statements, paintings, maps) and conclude that data-structures are not really internal representations at all. But if one takes the latter line, the modest successes of AI simply serve

to undercut our first premise: it is no longer obvious that psychology needs internal representations; internal pseudo-representations may do just as well.

It is certainly tempting to argue that since AI has provided us with the only known way of solving Hume's Problem, albeit for very restrictive systems, it must be on the right track, and its categories must be psychologically real, but one might well be falling into Danger (2) if one did. We can all be relieved and encouraged to learn that there is *a* way of solving Hume's Problem, but it has yet to be shown that AI's way is the only way it can be done.

AI has made a major contribution to philosophy and psychology by revealing a particular way in which simple cases of Hume's Problem can be solved. What else has it accomplished of interest to philosophers? I will close by just drawing attention to the two main areas where I think the AI approach is of particular relevance to philosophy.

For many years philosophers and psychologists have debated (with scant interdisciplinary communication) about the existence and nature of mental images. These discussions have been relatively fruitless, largely, I think, because neither side had any idea of how to come to grips with Hume's Problem. Recent work in AI, however, has recast the issues in a clearly more perspicuous and powerful framework, and anyone hoping to resolve this ancient issue will find help in the AI discussions.<sup>15</sup>

The second main area of philosophical interest, in my view, is the so-called "frame problem."<sup>16</sup> The frame problem is an abstract *epistemological* problem that was in effect discovered by AI thought-experimentation. When a cognitive creature, an entity with many beliefs about the world, performs an act, the world changes and many of the creature's beliefs must be revised or updated. How? It cannot be that we perceive and notice *all* the changes (for one thing, many of the changes we *know* to occur do not occur in our perceptual fields), and hence it cannot be that we rely entirely on perceptual input to revise our beliefs. So we must have internal ways of up-dating our beliefs that will fill in the gaps and keep our internal model, the totality of our beliefs, roughly faithful to the world.

If one supposes, as philosophers traditionally have, that one's beliefs are a set of propositions, and reasoning is inference or deduction from members of the set, one is in for trouble, for it is quite clear (though still controversial) that systems relying only on such processes get swamped by combinatorial explosions in the updating effort. It seems that our entire conception of belief and reasoning must be radically revised if we are to explain the undeniable capacity of

human beings to keep their beliefs roughly consonant with the reality they live in.

I think one can find an *appreciation* of the frame problem in Kant (we *might* call the frame problem Kant's Problem) but unless one disciplines one's thought-experiments in the AI manner, philosophical proposals of solutions to the problem, including Kant's of course, can be viewed as at best suggestive, at worst mere wishful thinking.

I do not want to suggest that philosophers abandon traditional philosophical methods and retrain themselves as AI workers. There is plenty of work to do by thought-experimentation and argumentation, disciplined by the canons of philosophical method and informed by the philosophical tradition. Some of the most influential recent work in AI (e.g., Minsky's papers on "Frames") is loaded with recognizably philosophical speculations of a relatively unsophisticated nature. Philosophers, I have said, should study AI. Should AI workers study philosophy? Yes, unless they are content to reinvent the wheel every few days. When AI reinvents a wheel, it is typically square, or at best hexagonal, and can only make a few hundred revolutions before it stops. Philosopher's wheels, on the other hand, are perfect circles, require *in principle* no lubrication, and can go in at least two directions at once. Clearly a meeting of minds is in order.\*

\*I am indebted to Margaret Boden for valuable advice on an early draft of this paper. Her *Artificial Intelligence and Natural Man* (Harvester, 1977), provides an excellent introduction to the field of AI for philosophers.



III

Objects of Consciousness and the  
Nature of Experience



## Are Dreams Experiences?

The “received view” of dreams is that they are *experiences that occur during sleep*, experiences which we can often recall upon waking. Enlarged, the received view is that dreams consist of sensations, thoughts, impressions, and so forth, usually composed into coherent narratives or adventures, occurring somehow in awareness or consciousness, though in some other sense or way the dreamer is *unconscious* during the episode.\* *Received* it certainly is; as Norman Malcolm pointed out in his book, *Dreaming*, not only has it been virtually unchallenged, it has been explicitly endorsed by Aristotle, Descartes, Kant, Russell, Moore, and Freud.<sup>1</sup> That was in 1959, and I think it is fair to say that in spite of Malcolm’s arguments against the received view, it is still the received view. I want to reopen the case, and though my aims and presuppositions are quite antagonistic to Malcolm’s, those familiar with his attack will see many points at which my discussion agrees with and gains insight and direction from his. I will not, though, go into a detailed extraction and defense of what I find valuable in Malcolm’s book. My immediate purpose in what follows is to undermine the authority of the received view of dreams. My larger purpose is to introduce a view about the relationship between experience and memory that I plan to incorporate into a physicalistic theory of consciousness, a theory considerably different from the theory I have hitherto defended.<sup>2</sup>

The most scandalous conclusion that Malcolm attempted to draw

\*Cf. Hilary Putnam’s version of “a natural lexical definition”: “a series of impressions (visual, etc.) occurring during sleep; usually appearing to the subject to be of people, objects, etc.; frequently remembered upon awakening” (“Dreaming and ‘Depth Grammar’” in R. J. Butler, ed., *Analytical Philosophy* (Oxford, 1962):224.

from his analysis of the concept of dreaming was to the effect that contemporary dream research by psychologists and other scientists was conceptually confused, misguided, ultimately simply *irrelevant* to dreaming.<sup>3</sup> This conclusion strikes many as bizarre and impertinent. If scientists can study waking experience, waking sensation, thought, imagination, consciousness, they can surely study the varieties of these phenomena that occur during sleep, in dreaming. This riposte is not, of course, a consideration that would impress Malcolm, for it is simply an announcement of faith in the received view, the view that dreams do consist of sensations, thoughts, and so forth occurring during sleep, and Malcolm already knows that the view he is attacking inspires such faith. In any event, as everyone expected, Malcolm's words have had little or no discouraging effect on dream researchers. Their work continues apace to this day, apparently with a degree of fruition that makes a mockery of Malcolm's view. So let us suppose, *contra* Malcolm, that the researchers are neither the perpetrators nor the victims of a conceptual crime, and see where it leads us. Let us suppose that the dream researcher's concept of dreaming is not only received, but the true and unconfused concept of dreaming. What are the prospects, then, for the scientific elaboration of the received view?

It is well known that periods of rapid eye movements (REMs) occur during sleep, and correlate well with subsequent reports of having dreamed. There are also characteristic EEG patterns usually concurrent with the REM episodes, and other physiological correlates that go to suggest that dreams do indeed occur during sleep, and can now be timed, confirmed to occur, and measured in all manner of ways. One tantalizing finding has been the apparent occasional content-relativity of the REMs. A person whose REMs are predominantly horizontal is awakened and reports a dream in which he watched two people throwing tomatoes at each other. A predominantly vertical pattern in REMs is correlated with a dream report of picking basketballs off the floor and throwing them up at the basket.<sup>4</sup> A neurophysiological model\* of dreaming would plausibly construe these REMs as relatively gross and peripheral effects of a more determinate content-relative process deeper in the brain, which we might hope some day to *translate*, in this sense: we might be able to *predict* from certain physiological events observed during sleep that the subsequent dream report would

\*Putnam (*op. cit.*) points out that a crucial lacuna in Malcolm's verificationist arguments against REMs as evidence confirming the received view is his failure to consider the confirmation relations arising from the use of developed theories and models (p. 226). At a number of points this paper attempts to fill that gap.

allude to, for example, fear, falling from a height, eating something cold, even (in the Golden Age of neurocryptography) buying a train ticket to New Haven for \$12.65 and then forgetting which pocket it was in. The prospect of a *generalized* capacity to predict dream narratives in such detail would be vanishingly small in the absence of a highly systematic and well-entrenched theory of representation in the brain, but let us suppose for the nonce that such a theory is not only in principle possible, but the natural culmination of the research strategies that are already achieving modest success in “translating” relatively gross and peripheral nervous-system activity.\*

Now some people claim never to dream, and many people waken to report that they have dreamed but cannot recall any details. The latter usually have a strong conviction that the dream *did* have details, though they cannot recall them, and even when we can recall our dreams, the memories fade very fast, and the mere act of expressing them seems to interfere, to speed up the memory loss. Here the impression of details *there then* but *now lost* is very strong indeed. REM researchers now confidently state that their research shows that *everybody* has dreams (and every night); some of us just seldom—or never—recall them. It must be unsettling to be assured that one has dreamed when one is positive one has not; Malcolm could be expected to diagnose one’s reaction to such an assurance as the shudder of conceptual violation,<sup>5</sup> but that would be an overstatement. The data of common experience strongly suggest a gradation in people’s capacities to recall (both dreams and other items), and it should be nothing worse than an odd but obvious implication of the received view that one could *dream* without recalling, just as one can promise without recalling, or be raucously drunk without recalling.

Guided by common experience and the received view, then, we can imagine our scientists of the future isolating the memory mechanisms responsible for dream recall, and finding ways of chemically facilitating or inhibiting them. This is surely plausible; research into the chemistry of memory already suggests which chemicals might have these powers. We would expect that the scientists’ claims to a theory of the dream-recall mechanism would be buttressed by systematic ties to a theory of memory mechanisms in general and by results, such as, perhaps, their ability to cure the dream-amnesiac.

\*I have in mind such work as Hubel and Wiesel’s “translation” of optic nerve signals in the cat. I argue against optimism regarding the prospects for a generalized neural theory of representation in “Brain Writing and Mind Reading” (Chapter 3 of this volume).

So we imagine future dream theory to posit two largely separable processes: first, there are neural events during sleep (more specifically during REM periods having certain characteristic EEG correlates) that systematically represent (are systematically correlatable with) the “events occurring in the dream”, and during this process there is a second, memory-loading process so that these events can be recalled on waking (when the memory process works). Dreams are *presented*, and simultaneously *recorded* in memory, and we might be able to interfere with or prevent the recording without disturbing the presentation.

This posited process of memory-loading and playback must be saved from simplistic interpretation if we are to maintain any vestige of realism for our fantasy. It is rarely if ever the case that a dreamer awakens and proceeds to recite with vacant stare a fixed narrative. Dream recall is like recall generally. We interpret, extrapolate, revise; it sometimes seems that we “relive” the incidents and *draw conclusions* from this reliving—conclusions that are then expressed in what we actually *compose* then and there as our recollections. It is not easy to analyze what must be going on when this happens. What is the *raw material*, the evidence, the basis for these reconstructions we call recollections?

Consider a fictional example. John Dean, a recently acclaimed virtuoso of recollection, is asked about a certain meeting in the Oval Office. Was Haldeman present? Consider some possible replies.

- (1) “No.”
- (2) “I can’t (or don’t) recall his being there.”
- (3) “I distinctly recall that he was not there.”
- (4) “I remember noticing (remarking) at the time that he was not there.”

If Dean says (1) we will suspect that he is saying less than he *can* say, even if what he says is sincere and even true. At the other extreme, (4) seems to be a nearly *complete* report of the relevant part of Dean’s memory. Answer (2), unlike all the others, reports an inability, a blank. Under the right circumstances, though, it carries about as strong a pragmatic implication of Haldeman’s absence as any of the others (we ask: could Dean conceivably fail to recall Haldeman’s presence if Haldeman had been there?). The stronger these pragmatic implications, the more disingenuous an answer like (2) will seem. Consider: “Was Dan Rather at that meeting in the Oval Office?” “I can’t *recall* his being there.” The answer is seen to be disingenuous because we know Dean knows, and we know, the additional supporting premises which,

in conjunction with (2), imply something like (1), and we expect Dean to be reasonable and draw this conclusion for—and with—us. Then what should Dean say, if asked the question about Dan Rather? Certainly not (4), unless the paranoia in the White House in those days knew no bounds, but (3) can be heard to carry a similar, if weaker, implication. We would not expect Dean to say this because it suggests (presumably misleadingly) that his answer is closer to being *given* in his recollection, less a conclusion quickly drawn. (1) is clearly the best answer on the list under these circumstances. It *looks* like a conclusion he reaches on the basis of things he remembers. He remembers Nixon and Ehrlichman talking with him, forming a sort of triangle in the room, and on the basis of *this* he concludes that Haldeman, and Rather, were absent, though he took no notice of the fact at the time, or if he did he has forgotten it. Now suppose Dean says (1). Perhaps when he does this he recalls in his mind this triangle, but does not bother to tell us that—he does not close his eyes on the witness stand and do a little phenomenology for us; he simply offers up his conclusion as a dictate of memory. But he need not have gone through this conscious process of reliving and reasoning at all. He may say, directly, “No,” and if he is pressed to be more forthcoming, any reasoning he offers based on other things he recalls will not be expressing any reasoning he knows he went through before his initial negative reply. He may not even be able to explain why or how his memory dictates this answer to the question, and yet be sure, and deservedly sure, that his reply is a sincere and reliable dictate of memory.

To summarize: sometimes we can sincerely answer a question of recollection with an answer like (4), but often we cannot, and sometimes we draw a blank, but in *all* these cases there are conclusions we can draw based on what in some sense we directly remember in conjunction with common and proprietary knowledge, and these conclusions need not be drawn in a process of *conscious* reasoning. Whatever it is that is directly remembered can play its evidentiary role in prompting an answer of recollection without coming into consciousness. This suggests that when we remember some event, there is some limited amount of information that is *there*, not necessarily in consciousness but available in one way or another for utilization in composing our recollections and answering questions we or others raise. Perhaps what occupies this functional position is an immensely detailed recording of our experience to which our later access in normally imperfect and partial (although under hypnosis it may improve). Perhaps there is enough information in this position to reconstitute completely our past experience and present us, under special circumstances, with a vivid

hallucination of reliving the event.\* However much is in this position in Dean, however, it is not possible that Dan Rather's absence is there except by implication, for his absence was not experienced by Dean at the time, any more than up to this moment you have been experiencing Rather's absence from this room. What the posited memory-loading process records, then, is whatever occupies this functional position at a later time. The "playback" of dream recollections, like other recollections, is presumably seldom if ever complete or uninterpreted, and often bits of information are utilized in making memory claims without being played back in consciousness at all.

In dreaming there is also a third process that is distinguished both in the layman's version of the received view and in fancier theories, and that is the *composition* of what is presented and recorded. In various ways this process exhibits intelligence: dream stories are usually coherent and realistic (even surrealism has a realistic background), and are often gripping, complex, and of course loaded with symbolism. Dream composition utilizes the dreamer's general and particular knowledge, her recent and distant experience, and is guided in familiar ways by her fears and desires, covert and overt.

Studying these three processes will require tampering with them, and we can imagine that the researchers will acquire the technological virtuosity to be able to influence, direct, or alter the composition process, to stop, restart, or even transpose the presentation process as it occurs, to prevent or distort the recording process. We can even imagine that they will be able to obliterate the "veridical" dream memory and substitute for it an undreamed narrative. This eventuality would produce a strange result indeed. Our dreamer would wake up and report her dream, only to be assured by the researcher that she never dreamed *that* dream, but rather another, which they proceed to relate to her. Malcolm sees that the scientific elaboration of the received view countenances such a possibility-in-principle and for him this amounts to a *reductio ad absurdum* of the received view,<sup>6</sup> but again, this is an overreaction to an admittedly strange circumstance. Given the state of the art of dream research today, were someone to contradict my clear recollection of what I had just dreamed, my utter skepticism would be warranted, but the science-fictional situation envisaged would be quite different. Not only would the researchers have proved their powers by correctly predicting dream recollections on numerous occasions, but they would have a theory that explained their successes.

\*Cf. Wilder Penfield's descriptions of electrode-induced memory hallucinations, in *The Excitable Cortex in Conscious Man*, (Springfield, Illinois, 1958).



And we need not suppose the dream they related to the dreamer would be entirely *alien* to her ears, even though she had no recollection of it (and in fact a competing recollection). Suppose it recounted an adventure with some secretly loved acquaintance of hers, a person unknown to the researchers. The stone wall of skepticism would begin to crumble.

The story told so far does not, I take it, exhibit the conceptual chaos Malcolm imagines; strange as it is, I do not think it would evoke in the layman, our custodian of ordinary concepts, the nausea of incomprehension. As a premise for a science-fiction novel it would be almost pedestrian in its lack of conceptual horizon-bending.

But perhaps this is not at all the way the theory of dreaming will develop. Malcolm notes in passing that it has been suggested by some researchers that dreams may occur during the moments of waking, not during the prior REM periods. Why would anyone conjecture this? Perhaps you have had a dream leading logically and coherently up to a climax in which you are shot, whereupon you wake up and are told that a truck has just backfired outside your open window. Or you are fleeing someone in a building, you climb out a window, walk along the ledge, then fall—and wake up on the floor having fallen out of bed. In a recent dream of mine I searched long and far for a neighbor's goat; when at last I found her she bleated *baa-a-a*—and I awoke to find her bleat merging perfectly with the buzz of an electric alarm clock I had not used or heard for months. Many people, I find, have anecdotes like this to relate, but the scientific literature disparages them, and I can find only one remotely well-documented case from an experiment: different stimuli were being used to waken dreamers, and one subject was awakened by dripping cold water on his back. He related a dream in which he was singing in an opera. Suddenly he *heard and saw* that the soprano had been struck by debris falling from above; he ran to her and as he bent over her, felt water dripping on his back.<sup>7</sup>

What are we to make of these reports? The elaboration of the received view we have just sketched can deal with them, but at a high cost: precognition. If the terminal events in these dreams are strongly *prepared for* by the narrative, if they do not consist of radically juxtaposed turns in the narrative (for example, the goat turns into a telephone and starts ringing), then the composition process must have been directed by something having “knowledge” of the future. That is too high a price for most of us to pay, no doubt. Perhaps all these anecdotes succumb to a mixture of reasonable skepticism, statistics (coincidences do happen, and are to be “expected” once in a blue moon), the discovery of subtle influences from the environment, and

various other deflating redescrptions. But if all else failed we could devise any number of variant dream theories that accommodated these "miracles" in less than miraculous ways. Perhaps, to echo the earlier conjecture, dreams are composed and presented *very fast* in the interval between bang, bump, or buzz and full consciousness, with some short delay system postponing the full "perception" of the noise in the dream until the presentation of the narrative is ready for it. Or perhaps in that short interval dreams are composed, presented, and recorded *backwards*, and then remembered front to back. Or perhaps there is a "library" in the brain of undreamed dreams with various indexed endings, and the bang or bump or buzz has the effect of retrieving an appropriate dream and inserting it, cassette-like, in the memory mechanism.

None of these theories can be viewed as a mere variation or rival elaboration of the received view. If one of them is true, then the received view is false. And since these rival theories, including the theory inspired by the received view, are all empirical, subject to confirmation and refutation, and since the rival theories even have some (admittedly anecdotal) evidence in their favor, we are constrained to admit that the received view might simply turn out to be false: dreams, it might turn out, are not what we took them to be—or perhaps we would say that it turns out that there are no dreams after all, only dream "recollections" produced in the manner described in our confirmed theory, whichever it is. Malcolm sees that all this is implied by the received view, and takes it to be yet another *reductio ad absurdum* of it: any view that could permit the discovery that "we are always only under the *illusion* of having had a dream" is "senseless".<sup>8</sup> But again, Malcolm's response to this implication is too drastic. The claim that we had been fooled for millennia into believing in dreams would be hard to swallow, but then we would not have to swallow it unless it had the backing of a strongly confirmed scientific theory, and then this claim would put no greater strain on our credulity than we have already endured from the claims of Copernicus, Einstein, and others. It would be rather like learning that dream-recall was like *déjà vu*—it only *seemed* that you had experienced it before—and once you believed *that*, it would no longer even seem (as strongly) that you were recalling. The experience of "dream recall" would change for us.\*

My attack on the received view is not, however, a straightforward

\*Cf. Putnam, *op. cit.*, p. 227. The naïve subject of *déjà vu* says, "I vaguely remember experiencing all this before"; the sophisticated subject is not even tempted to say this, but says, perhaps, "Hm, I'm having a *déjà vu* experience right now." The experience has changed.

empirical attack. I do not wish to aver that anecdotal evidence about dream anticipation disproves the received view, but I do want to consider in more detail what the issues would be were a rival to the received view to gain support. I hope to show that the received view is more vulnerable to empirical disconfirmation than its status as the received view would lead us to expect. Of the rival theories, the cassette-library theory runs most strongly against our pretheoretical convictions, for on the other two there still is some vestige of the presumed presentation process: it is just much faster than we had expected, or happens backwards. On the cassette view, our "precognitive" dreams are never dreamed at all, but just spuriously "recalled" on waking. If our memory mechanisms were empty until the moment of waking, and then received a whole precomposed dream narrative in one lump, the idea that precognitive dreams are *experienced episodes* during sleep would have to go by the board.

Suppose we generalize the cassette theory to cover all dreams: all dream narratives are composed directly into memory banks; which, if any, of these is available to waking recollection depends on various factors—precedence of composition, topicality of waking stimulus, degree of "repression", and so forth. On this view, the process of presentation has vanished, and although the dream cassettes would have to be filled at some time by a composition process, that process might well occur during our waking hours, and spread over months (it takes a long time to write a good story). The composition might even have occurred aeons before our birth; we might have an *innate* library of undreamed dream cassettes ready for appropriate insertion in the playback mechanism. Stranger things have been claimed. Even on the received view, the composition process is an unconscious or subconscious process of which we normally have no more *experience* than of the processes regulating our metabolism; otherwise dreams could not be suspenseful. (I say "normally" for there does seem to be the phenomenon of self-conscious dreaming, where we tinker with a dream, run it by several times, attempt to resume it where it left off. Here the theatrical metaphor that enlivens the received view seems particularly apt. After tinkering like the playwright, we must sit back, get ourselves back into the audience mood, suspend disbelief, and re-enter the play. Some researchers call these occasions *lucid dreams*. But usually we are not privy to the composition process at all, and so have no inkling about when it might occur.) Research might give us good grounds for believing that dream narratives that were composed onto cassettes in the morning decayed faster than cassettes composed in the afternoon, or during meals.

A more likely finding of the cassette-theorist would be that the composition process occurs during sleep, and more particularly, during periods of rapid eye movements, with characteristic EEG patterns. One might even be able to “translate” the composition process—that is, predict dream recollections from data about the composition process. This theory looks suspiciously like the elaboration of the received theory, except that it lacks the presentation process. Cassette narratives, we are told, are composed in narrative order, and long narratives take longer to compose, and the decay time for cassettes in storage is usually quite short; normally the dream one “recalls” on waking was composed just minutes earlier, a fact attested to by the occasional cases of content-relativity in one of the by-products of cassette composition: rapid eye movements. On this theory dream memories are produced just the way the received theory says they are, except for one crucial thing: the process of dream-memory production is entirely unconscious, involves no awareness or experiencing at all. Even “lucid dreams” can be accommodated easily on this hypothesis, as follows: although the composition and recording processes are entirely unconscious, on occasion the composition process inserts traces of itself into the recording via the literary conceit of a dream within a dream.\*

Now we have a challenge to the received view worth reckoning with. It apparently accounts for all the data of the REM researchers as well as the received view does, so there is no reason for sober investigators not to adopt the cassette theory forthwith if it has any advantages over the received view. And it seems that it does: it has a simple explanation of precognitive dreams (if there are any) and it posits one less process by eliminating a presentation process whose point begins to be lost.

But what greater point could a process have? In its presence we have experience; in its absence we have none. As Thomas Nagel would put it, the central issue between these two theories appears to be whether or not it is like anything to dream.<sup>9</sup> On the cassette theory it is not like anything to dream, although it is like something *to have dreamed*. On the cassette theory, dreams are not experiences we have during sleep; where we had thought there were dreams, there is only an unconscious composition process and an equally unconscious memory-loading process.

A few years ago there was a flurry of experimentation in learning-while-you-sleep. Tape recordings of textbooks were played in the

\*For more on lucid dreams and their accommodation by the cassette theory, see Kathleen Emmett, “Oneiric Experiences” and my “The Onus Re Experiences: a Reply to Emmett”, both forthcoming in *Philosophical Studies*.

sleeper's room, and tests were run to see if there were any subsequent signs of learning. As I recall, the results were negative, but some people thought the results were positive. If you had asked one of them *what it was like* to learn in one's sleep, the reply would presumably have been: "It was not like anything at all—I was sound asleep at the time. I went to sleep not knowing any geography and woke up knowing quite a bit, but don't ask me what it was like. It wasn't like anything." If the cassette theory of dreams is true, dream-recollection production is a similarly unexperienced process. If asked what it is like to dream one *ought* to say (because it would be the truth): "It is not like anything. I go to sleep and when I wake up I find I have a tale to tell, a 'recollection' as it were." It is Malcolm's view that this is what we ought to say, but Malcolm is not an explicit champion of the cassette theory or any other empirical theory of dreaming. His reasons, as we shall see, are derived from "conceptual analysis". But whatever the reasons are, the conclusion seems outrageous. *We all know better*, we think. But do we? We are faced with two strikingly different positions about what happens when we dream, and one of these, the received view, we are not just loath to give up; we find it virtually unintelligible that we could be wrong about it. And yet the point of difference between it (as elaborated into a theory by scientists) and its rival, the cassette theory, is apparently a technical, theoretical matter about which the layman's biases, his everyday experience, and even his personal recollections of dreams are without authority or even weight. What should we do? Sit back and wait for the experts to tell us, hoping against hope that dreams will turn out to be, after all, experiences? That seems ridiculous.

If that seems ridiculous, perhaps it *is* ridiculous. Can some way be found to protect the received view from the possibility of losing this contest? If we do not for a minute believe it could lose, we must suppose there is some principled explanation of this. One might set out in a verificationist manner.\* What could possibly settle the issue between the received view and the cassette theory if subjects' recollections were deemed neutral? The conclusion of one view is that dreams are experiences, and of the other, that they are not; but if subjects' recollections were not held to be *critical*, nothing else could count as evidence for or against the rival theories, at least with regard to this disputed conclusion. Therefore the claimed difference between the two theories is illusory, or perhaps we should say they are both pseudo-theories. This

\*This argument is inspired by the verificationist arguments of Malcolm, and its rebuttal is inspired by Putnam's objections, but Malcolm does not commit himself to this argument.

will not do. We can easily imagine the two theories to share a concept of experience, and even to agree on which data would go to show that dreams were, in this shared technical sense, experiences. Nor would this technical concept of experience have to look all that unordinary. We have many common ways of distinguishing which among the events that impinge on us are experienced and which are not, and we can imagine these theories to build from these ordinary distinctions a powerful shared set of well-confirmed empirically necessary and sufficient conditions for events to be experienced. If, for instance, some part of the brain is invariably active in some characteristic way when some event in waking life is, as we ordinarily say, experienced, and if moreover we have a theory that says why this should be so, the absence of such brain activity during REM periods would look bad for the received view and good for the cassette view.

But if that is what we should look for, the received view is in trouble, for one routinely recognized condition for having an experience is that one be conscious, or awake, and dreamers are not. A well-confirmed physiological condition for this is that one's reticular activating system be "on", which it is not during sleep. The fact that one is in a sound sleep goes a long way to confirming that one is *not* having experiences, as *ordinarily* understood. Malcolm would make this criterial, but that is one more overstatement. Lack of reticular system activity strongly suggests that nothing is being experienced during REMs, but the defender of the received view can plausibly reply that reticular activation is only a condition of *normal* experience, and can point to the frequent occurrence during REM periods of the normal physiological accompaniments of fear, anxiety, delight, and arousal as considerations in favor of an extended concept of experience. How could one exhibit an emotional reaction to something not even experienced? The debate would not stop there, but we need not follow it further now. The fact remains that the physiological data would be clearly relevant evidence in the dispute between the theories, and not all the evidence is on the side of the received view.

Still, one might say, the very relevance of physiological evidence shows the dispute not to involve our ordinary concept of experience at all, but only a technical substitute. For suppose we were told without further elaboration that the theory inspired by the received view had won the debate, had proved to be the better theory. We would not know what, if anything, had been confirmed by this finding. Which of our hunches and biases would be thereby vindicated, and are any of them truly in jeopardy?

This plausible rhetorical question suggests that none of our precious

preconceptions about dreaming *could* be in jeopardy, a conclusion that “conceptual analysis” might discover for us. How might this be done? Let us return to the comparison between the cassette view of dreams and the speculation that one might learn in one’s sleep. I suggested that subjects in either circumstance should, on waking, deny that it was like anything to have undergone the phenomenon. But there would be a crucial difference in their waking states, presumably. For the dreamer, unlike the sleep-learner, would probably want to add to his disclaimer: “Of course it *seems to me* to have been like something!” The sleep learner has new knowledge, or new beliefs, but not new *memories*. This is surely an important difference, but just what difference does it make? Is it that the claim:

(5) It was not like anything, but it seems to me to have been like something,

is a covert contradiction? Can one sustain the following principle?

(6) If it seems to have been like something, it was like something.

The present tense version of the principle is unassailable:

(7) If it seems to me to be like something to be *x*, then it is like something to me to be *x*.

That is what we mean when we talk of what it is like: how it seems to us.\* When we try to make the principle extend through memory to the past, however, we run into difficulties. There is no good reason to deny that memories can be spurious, and there is plenty of confirmation that they can. This is somewhat obscured by some looseness in our understanding of the verb “remember”. Sometimes we draw a distinction between remembering and seeming to remember such that remembering, like knowing, is veridical. On this reading it follows that if you remember something to have been *x*, it was *x*. If it was not *x*, you only seem to remember that it was. But when I say, about a restaurant we are dining in, “This isn’t the way I remember it,” my claim is equivocal. I may not be claiming the restaurant has changed—it may be that my memory is at fault. On this reading of “remember” there is still a distinction between remembering and seeming to remember, but it is not a distinction with veridicality on one side: for example one tells a tale of one’s childhood that is shown to be false and one wonders whether one has mistaken fantasizing or confabulating for (mis)remembering. On *either* reading, however, there is no claim that can be made of the form:

\*Cf. Nagel, “What Is It Like to Be a Bat?”:440n. “[T]he analogical form of the English expression ‘what is it *like*’ is misleading. It does not mean ‘what (in our experience) it resembles’, but rather ‘how it is for the subject himself’.”

(8) Since I remember it to have been like something, it was like something.

On the first reading of "remember", the claim, while logically impeccable, does no work unless one claims a capacity to tell one's memories from one's seeming memories that one simply does not have. On the second reading, even if we could always tell fantasy from memory the consequent would not follow. So (5) represents a possible state of affairs. We had in fact already countenanced this state of affairs as an abnormality in supposing that the dream researchers could, by tampering, insert a spurious dream recollection. Now we are countenancing it as a possible and not even improbable account of the normal case.

Malcolm sees that nothing like (6) or (8) can be exploited in this context; we can seem to have had an experience when we have not, and for just this reason he denies that dreams are experiences! His argument is that *since* one can be under the impression that one has had an experience and yet not have had it, and since if one is under the impression that one has had a dream, one *has* had a dream,\* having had a dream cannot be having had an experience; hence, dreams are not experiences.

This "criteriological" move has a curious consequence: it "saves" the authority of the wakened dream-recaller, and this *looks like* a rescue of subjectivity from the clutches of objective science, but it "saves" dreaming only at the expense of experience. What Malcolm sees is that if we permit a distinction between "remembering" and "seeming to remember" to apply to dream recollections, the concept of dreaming is cast adrift from any criterial anchoring to first-person reports, and becomes (or is revealed to be) a theoretical concept. Once we grant that subjective, introspective or retrospective evidence does not have the authority to settle questions about the nature of dreams—for instance, whether dreams are experiences—we have to turn to the other data, the behavior and physiology of dreamers, and to the relative strengths of the theories of these, if we are to settle the question, a question which the subject is not in a privileged position to answer.

\*"That he really had a dream and that he is under the impression that he had a dream: these are the same thing" ("Dreaming and Skepticism",:32). This is the central premise of Malcolm's work on dreaming, and one he gets from Wittgenstein: "The question whether the dreamer's memory deceives him when he reports the dream after waking cannot arise unless indeed we introduce a completely new criterion for the report's 'agreeing' with the dream, a criterion which gives us a concept of 'truth' as distinguished from 'truthfulness' here" (*Philosophical Investigations*:222-223). It is Malcolm's unswerving loyalty to this remark that forces his account into such notorious claims.



Malcolm avoids this by denying that dreams are experiences, but this only concedes that one does *not* have a privileged opinion about one's own past experiences.\* This concession is unavoidable, I think, and Malcolm's is not the only philosophic position caused embarrassment by it. A defender of the subjective realm such as Nagel must grant that in general, whether or not it was like something to be *x*, whether or not the subject *experienced* being *x*—questions that *define* the subjective realm—are questions about which the subject's subsequent subjective opinion is not authoritative. But if the subject's own convictions do not settle the matter, and if, as Nagel holds, no objective considerations are conclusive either, the subjective realm floats out of ken altogether, except perhaps for the subject's convictions about the specious present. Dreams are particularly vulnerable in this regard only because, as Malcolm observes, sleepers do not and cannot *express* current convictions about the specious present (if they have any) while they are dreaming. Since our only expressible access to dreams is retrospective, dreams are particularly vulnerable, but they are not alone. The argument we have been considering is more general; the dispute between the rival theories of memory-loading can be extended beyond dreaming to all experience. For instance, just now, while you were reading my remarks about Nagel, were you experiencing the peripheral sights and sounds available in your environment? Of course you were, you say, and you can prove it to your own complete satisfaction by closing your eyes and recalling a variety of events or conditions that co-occurred with your reading those remarks. While not *central* in your consciousness at the time, they were certainly *there*, being *experienced*, as your recollections show. But the cassette theorist, emboldened by the success with dreams, puts forward the *subliminal peripheral recollection-production* theory, the view that the variety of peripheral details in such cases are not consciously experienced, but merely unconsciously recorded for subsequent recall. Events outside our immediate attention are not experienced at all, our theorist says, but they do have subliminal effects on short-term memory. Our capacity to recall them for a short period does not establish that they were experienced, any more than our capacity to "recall" dreams shows that they were experienced. But this is nonsense, you say: *recording those peripheral items for subsequent recollection just is experiencing them.*

If only this bold claim were true! Look what it would do for us. The

\*Sometimes Malcolm seems to want to "save" all "private states" in this way, thus either having to deny that experiences are private states, or having to adopt after all some principle like (8). See *Dreaming*, p. 55.

difference between the received view of dreams and the cassette theory would collapse; the presumably unconscious memory-loading process of the cassette theory would turn out to be the very presentation process dear to the received view. A "conceptual relationship" could be established between experience and memory that avoided the difficulties heretofore encountered in such claims, as follows: The conceptual relationship is *not* between experiencing and subsequent subjective convictions of memory (the latter are *not* criterial), but between experiencing and something perfectly objective: the laying down thereupon of a memory trace—for however short a time and regardless of subsequent success or failure at recollection.\* The conceptual relationship would be identity. *Experire est recordare*.

Much can be said in support of this principle, but at this time I will restrict myself to a few brief persuasions. First, is remembering a *necessary* condition for experiencing? Arguably, yes, if you grant that memories may not last long. The idea of a subject, an "I", experiencing each successive state in a stream of consciousness with *no* recollection of its predecessors, is a hopelessly impoverished model of experience and experiencers. The *familiarity* and *continuity* in the world of current experiences is a necessary background for recognition and discrimination, and only short-term memory can provide this. Items that come and go so fast, or so inconspicuously, as to leave no reverberations behind in memory at all, are plausibly viewed as simply not experienced. So if remembering is a necessary condition, is it also a *sufficient* condition for experiencing? Yost and Kalish say so, without supporting argument: "Dreaming is a real experience. And since dreams can be remembered they must be conscious experiences."<sup>10</sup> Martin and Deutscher, in their article "Remembering", concur:

So long as we hold some sort of 'storage' or 'trace' account of memory, it follows that we can remember only what we have experienced, for it is in our experience of events that they "enter" the storehouse.<sup>11</sup>

So remembering, in the sense of storing away in the memory for some time, is arguably a necessary and sufficient condition for experiencing. These are, I think, philosophically respectable arguments for the claimed identity, and to them can be added an ulterior

\*Not completely regardless of subsequent success or failure at recollection, for identifying some process as the laying down of a memory trace is identifying some process by its function, and nothing that did not have as its normal effect enabling the subject to report truly about the past could be picked out functionally as the memory-loading process.

consideration which will appeal to physicalists if not to others. The proposed identity of experiencing and recording promises a striking simplification for physicalist theories of mind. The problematic (largely because utterly vague) presentation process vanishes as an extra phenomenon to be accounted for, and with it goes the even more mysterious *audience* or *recipient* of those presentations. In its place is just a relatively prosaic short-term memory capacity, the sort of thing for which rudimentary but suggestive physical models already abound.

The principle as it stands, however, is too strong, on two counts. Consider again Martin and Deutscher's commentary on the "storehouse" model of memory: "It is in our experience of events that they 'enter' the storehouse." What, though, of forcible or illegal entry? We need an account of something like *normal* entry into memory so that we can rule out, as experiences, such abnormally entered items as the undreamed dream surgically inserted by the dream researchers. We want to rule out such cases, not by declaring them impossible, for they are not, but by denying that they are experiences for the subject. As we shall see in a moment, the best way of doing this may have a surprising consequence. The second failing of our principle is simply that it lacks the status we have claimed for it. It is not self-evident; its denial is not a contradiction. We must not make the mistake of asserting that this is a discovered conceptual truth about experience and memory. We must understand it as a proposal, a theoretically promising adjustment in our ordinary concepts for which we may have to sacrifice some popular pre-conceptions. For instance, whether animals can be held to dream, or to experience anything, is rendered an uncertainty depending on what we mean by *recall*. Can animals *recall* events? If not, they cannot have experiences. More radically, subjective authority about experience goes by the board entirely. Still, we get a lot in return, not the least of which is a way of diagnosing and dismissing the Pickwickian hypothesis of subliminal peripheral recollection-production.

We are still not out of the woods on dreaming, though, for we must define normal memory-entry in such a way as to admit ordinary experience and exclude tampering and other odd cases.

When the memory gets loaded by accident or interference we will not want this to count as experience, and yet we want to grant that there is such a thing as nonveridical experience. The memory-loading that occurs during a hallucination occurs during abnormal circumstances, but not so abnormal as to lead us to deny that hallucinations are experiences. But look at a slightly different case (I do not know if this ever occurs, but it might). Suppose at noon Jones, who is wide awake, suffers some event in her brain that has a delayed effect: at

12:15 she will “recall” having seen a ghost at noon. Suppose her recollection is as vivid as you like, but suppose her actual behavior at noon (and up until recollection at 12:15) showed no trace of horror, surprise, or cognizance of anything untoward. Had she shown any signs at noon of being under the impression that something bizarre was happening, we would be strongly inclined to say she had had a hallucination then, was experiencing it then, even if she could not recount it to us until fifteen minutes later. But since she did not *react* in any such telling way at noon, but proceeded about her business, we are strongly inclined to say the hallucination occurred later, at 12:15, and was a *hallucination of recollection* of something she had never experienced, even though the cause of the hallucination occurred at noon. Since the events responsible for her later capacity to recall did not contribute to her behavior-controlling state at the time, they did not enter her experience then, whatever their later repercussions. But then when we apply this distinguishing principle to dreams, we find that it is quite likely that most dreams are not experiences. Whereas nightmares accompanied by moans, cries, cowering, and sweaty palms *would* be experiences, bad dreams dreamed in repose (though remembered in agony) would not be, unless, contrary to surface appearances, their entry into memory is accomplished by engagements of the whole behavior-controlling system sufficiently normal to distinguish these cases sharply from our imaginary delayed hallucination.\*

If it turns out that sleep, or at least that portion of sleep during which dreaming occurs, is a state of more or less peripheral paralysis or inactivity; if it turns out that most of the functional areas that are critical to the governance of our wide awake activity are in operation, then there will be good reason for drawing the lines around experience so that dreams are included. If not, there will be good reason to deny that dreams are experiences.

Some of the relevant data are already familiar. The occurrence of REMs suggests that more than a little of the visual processing system is active during dream periods, and it should be a fairly straightforward task—perhaps already accomplished—to determine just how much is. Even strongly positive results would not be overwhelming grounds for deciding that dreams are experiences, however, for in various sorts of hysteric or psychosomatic blindness there is substantial apparently normal activity in the visual processing system, and in so-called

\*Malcolm too sees an important distinction between “violent nightmares” and normal dreams dreamed in repose, a distinction that forces him to claim we have several different concepts of sleep. Only thus can he save as a *conceptual* truth the claim that we have no experiences while we sleep.

subliminal perception the same is true, and in neither case are we inclined to suppose that visual experience occurs. More compelling, in many ways, is the evidence that dreams serve a purpose: they seem to be used to redress emotional imbalances caused by frustrating experiences in waking life, to rationalize cognitive dissonances, allay anxieties, and so forth. When this task is too difficult, it seems, the dream mechanisms often go into a looping cycle; troubled people often report recurring obsessive dreams that haunt them night after night,\* so if a recurrent physiological process can be correlated with these dreams, it will appear to be a presentation process, and the presentation process will have a point: namely, to provide the emotional and cognitive-processing functional parts with the raw material for new syntheses, new accommodations, perhaps permitting a more stable or satisfying self-image for the dreamer. But even this function could easily be seen to be accomplished entirely *unconsciously*. The self-presentation tactics and perceptual interpretation ploys posited by theorists as diverse as Freud and Erving Goffman are no less plausible for being presumed to be entirely unconscious, and they serve a similar self-protective maintenance function. As Malcolm points out, dreamers' narratives can be used by Freudians and others as a valuable source of information about the internal processes that shape us, without our having to suppose that these are recollections of experiences.\*\*

It is an *open*, and *theoretical* question whether dreams fall inside or outside the boundary of experience.\*\*\* A plausible theory of experience will be one that does justice to *three* distinguishable families of intuitions we have about experience and consciousness: those dealing

\*I am indebted to Robert Nozick for raising this consideration.

\*\**Dreaming*, p. 122. Malcolm quotes with approval this methodological suggestion of Freud's from *A General Introduction to Psychoanalysis* (Garden City: Norton, 1943), p. 76.

Any disadvantage resulting from the uncertain recollection of dreams may be remedied by deciding that exactly what the dreamer tells is to count as the dream, and by ignoring all that he may have forgotten or altered in the process of recollection.

\*\*\*Foulkes, *The Psychology of Sleep* (New York: Scribners, 1966) cites a number of telling, if inconclusive, further observations: in one study no association was found between "the excitement value of dream content and heart or respiration rate" (p. 50), a datum to be balanced by the curious fact that there are usually action-potentials discoverable in the motor-neurons in the biceps of one who is asked to *imagine* bending one's arm; similar action-potentials are found in the arms of deaf mute dreamers—people who talk with their hands. There are also high levels of activity in the sensory cortex during dreaming sleep.

with the role of experience in guiding current behavior, those dealing with our *current* proclivities and capacities to say what we are experiencing, and those dealing with the *retrospective* or *recollective* capacity to say. In earlier work I have sharply distinguished the first and second of these, but underestimated the distinctness and importance of this third source of demands on a theory of consciousness. A theory that does justice to these distinct and often inharmonious demands must also do justice to a fourth: the functional salencies that emerge from empirical investigation. In the end, the concept of experience may not prove to differentiate any one thing of sufficient theoretical interest to warrant time spent in determining its boundaries. Were this to occur, the received view of dreams, like the lay view of experience in general, would not be so much disproved as rendered obsolete. It may seem inconceivable that this could happen, but arm-chair conceptual analysis is powerless to establish this.