JEROME BRUNER

ACTS OF MEANING

• CHAPTER ONE •

The Proper Study of Man

I WANT TO BEGIN with the Cognitive Revolution as my point of departure. That revolution was intended to bring "mind" back into the human sciences after a long cold winter of objectivism. But mine will not be the usual account of progress marching ever forward.[1] For, at least in my view, that revolution has now been diverted into issues that are marginal to the impulse that brought it into being. Indeed, it has been technicalized in a manner that even undermines that original impulse. This is not to say that it has failed: far from it, for cognitive science must surely be among the leading growth shares on the academic bourse. It may rather be that it has become diverted by success, a success whose technological virtuosity has cost dear. Some critics, perhaps unkindly, even argue that the new cognitive science, the child of the revolution, has gained its technical successes at the price of dehumanizing the very concept of mind it had sought to reestablish in psychology, and that it has thereby estranged much of psychology from the other human sciences and the humanities.[2]

I shall have more to say on these matters shortly. But before going on, let me give you the plan of this chapter and the ones that follow. Once our retrospective glance at the revolution is

done, I then want to turn directly to a preliminary exploration of a renewed cognitive revolution—a more interpretive approach to cognition concerned with "meaning-making," one that has been proliferating these last several years in anthropology, linguistics, philosophy, literary theory, psychology, and, it would almost seem, wherever one looks these days.[3] I rather suspect that this vigorous growth is an effort to recapture the original momentum of the first cognitive revolution. In later chapters, I shall try to fill in this preliminary sketch with some concrete illustration of research on the boundaries between psychology and its neighbors in the humanities and the social sciences, research that recaptures what I have called the originating impulse of the cognitive revolution.

Now let me tell you first what I and my friends thought the revolution was about back there in the late 1950s. It was, we thought, an all-out effort to establish meaning as the central concept of psychology—not stimuli and responses, not overtly observable behavior, not biological drives and their transformation, but meaning. It was *not* a revolution against behaviorism with the aim of transforming behaviorism into a better way of pursuing psychology by adding a little mentalism to it. Edward Tolman had done that, to little avail.[4] It was an altogether more profound revolution than that. Its aim was to discover and to describe formally the meanings that human beings created out of their encounters with the world, and then to propose hypotheses about what meaning-making processes were implicated. It focused upon the symbolic activities that human beings employed in constructing and in making sense not only of the world, but of themselves. Its aim was to prompt psychology to join forces with its sister interpretive disciplines in the humanities and in the social sciences. Indeed,

beneath the surface of the more computationally oriented cognitive science, this is precisely what has been happening—first slowly and now with increasing momentum. And so today one finds flourishing centers of cultural psychology, cognitive and interpretive anthropology, cognitive linguistics, and above all, a thriving worldwide enterprise that occupies itself as never before since Kant with the philosophy of mind and of language. It is probably a sign of the times that the two Jerusalem-Harvard Lecturers in the academic year 1989–90 represent, each in his own way, this very tradition—Professor Geertz in anthropology and myself in psychology.

The cognitive revolution as originally conceived virtually required that psychology join forces with anthropology and linguistics, philosophy and history, even with the discipline of law. It is no surprise and certainly not an accident that in those early years the advisory board of the Center for Cognitive Studies at Harvard included a philosopher, W. V. Quine, an intellectual historian, H. Stuart Hughes, and a linguist, Roman Jakobson. Or that among the Center's Fellows could be numbered almost as many philosophers, anthropologists, and linguists as there were proper psychologists—among them such exponents of the new constructivism as Nelson Goodman. As for the law, I must report that several distinguished members of that faculty came occasionally to our colloquia. One of them, Paul Freund, admitted he came because we at the Center, it seemed to him, were interested in how rules (like rules of grammar, rather than scientific laws) affected human action and that, after all, is what jurisprudence is about.[5]

I think it should be clear to you by now that we were not out to "reform" behaviorism, but to replace it. As my col-

league George Miller put it some years later, "We nailed our new credo to the door, and waited to see what would happen. All went very well, so well, in fact, that in the end we may have been the victims of our success."[6]

It would make an absorbing essay in the intellectual history of the last quarter-century to trace what happened to the originating impulse of the cognitive revolution, how it became fractionated and technicalized. The full story had best be left to the intellectual historians. All we need note now are a few signposts along the way, just enough of them to give a sense of the intellectual terrain on which we were all marching. Very early on, for example, emphasis began shifting from "meaning" to "information," from the *construction* of meaning to the *processing* of information. These are profoundly different matters. The key factor in the shift was the introduction of computation as the ruling metaphor and of computability as a necessary criterion of a good theoretical model. Information is indifferent with respect to meaning. In computational terms, information comprises an already precoded message in the system. Meaning is preassigned to messages. It is not an outcome of computation nor is it relevant to computation save in the arbitrary sense of assignment.

Information processing inscribes messages at or fetches them from an address in memory on instructions from a central control unit, or it holds them temporarily in a buffer store, and then manipulates them in prescribed ways: it lists, orders, combines, compares precoded information. The system that does all of these things is blind with respect to whether what is stored is words from Shakespeare's sonnets or numbers from a random number table. According to classic information theory, a message is informative if it reduces alternative choices.

This implies a code of established possible choices. The categories of possibility and the instances they comprise are processed according to the "syntax" of the system, its possible moves. Insofar as information in this dispensation can deal with meaning it is in the dictionary sense only: accessing stored lexical information according to a coded address. There are other "meaning-like" operations such as permuting a set of entries in order to test the resultants against a criterion, as in anagrams or Scrabble. But information processing cannot deal with anything beyond well-defined and arbitrary entries that can enter into specific relationships that are strictly governed by a program of elementary operations. Such a system cannot cope with vagueness, with polysemy, with metaphoric or connotative connections. When it seems to be doing so, it is a monkey in the British Museum, beating out the problem by a bone-crushing algorithm or taking a flyer on a risky heuristic. Information processing needs advance planning and precise rules.[7] It precludes such ill-formed questions as "How is the world organized in the mind of a Muslim fundamentalist?" or "How does the concept of Self differ in Homeric Greece and in the postindustrial world?" And it favors questions like "What is the optimum strategy for providing control information to an operator to ensure that a vehicle will be kept in a predetermined orbit?" We shall have much more to say later about meaning and the processes that create it. They are surprisingly remote from what is conventionally called "information processing."

It is not surprising, given that an Information Revolution was occurring throughout the postindustrial world, that such an emphasis should have developed. Psychology and the social sciences generally have always been sensitive, often oversensi-

tive, to the needs of the society that gives them shelter. And it has always been rather an intellectual reflex of academic psychology to redefine man and his mind in the light of new social requirements. Nor is it surprising that under such conditions interest should have shifted away, accordingly, from mind and meaning to computers and information. For computers and computational theory had by the early 1950s become the root metaphor for information processing. Given preestablished meaning categories well-formed enough within a domain to provide a basis for an operating code, a properly programmed computer could perform prodigies of information processing with a minimum set of operations, and that is technological heaven. Very soon, computing became the model of the mind, and in place of the concept of meaning there emerged the concept of computability. Cognitive processes were equated with the programs that could be run on a computational device, and the success of one's effort to "understand," say, memory or concept attainment, was one's ability realistically to simulate such human conceptualizing or human memorizing with a computer program.[8] This line of thinking was enormously aided by Turing's revolutionary insight that any computational program, no matter how complex, could be "imitated" by a much simpler Universal Turing Machine computing with a finite set of quite primitive operations. If one falls into the habit of thinking of those complex programs as "virtual minds" (to borrow Daniel Dennett's phrase), then it takes only a small but crucial step to go the whole way to believing that "real minds" and their processes, like "virtual" ones and theirs, could be "explained" in the same way.[9]

This new reductionism provided an astonishingly libertar-

ian program for the new cognitive science that was being born. It was so permissive, indeed, that even the old S-R learning theorist and associationist student of memory could come right back into the fold of the cognitive revolution so long as they wrapped their old concepts in the new terms of information processing. One did not have to truck with "mental" processes or with meaning at all. In place of stimuli and responses, there was input and output, with reinforcement laundered of its affective taint by being converted into a control element that fed information about the outcome of an operation back into the system. So long as there was a computable program, there was "mind."

At first this pun version of mind did not seem to provoke the traditional antimentalist panic among the seemingly converted behaviorists. In good time, though, new versions of old classically familiar controversies began to reemerge, particularly in connection with debates about the so-called architecture of cognition: whether it was to be conceived as a set of grammar-like hierarchically nesting rule structures for accepting, rejecting, or combining input, or whether, rather, it could be conceived of as a bottom-up connectionist network with completely distributed control as in the PDP (Parallel Distributed Processing) models, a model much like the old associationist doctrine, minus Herbart's creative synthesis. The first simulated the top-down, rationalist-mentalist tradition in psychology and moved easily back and forth between "real" minds and "virtual" ones; the second was a new version of what Gordon Allport used to mock in his lectures as "dust-bowl empiricism." East Coast computationalism dealt with such mindlike terms as rules, grammars, and the like. The West Coasters wanted no part of such simulated mentalism.

Soon, the battleground began looking increasingly traditional and familiar, though the vehicles that were racing over it had much more speed and much more formalistic horsepower. But whether their maneuvers had to do with the *mind* or only with the theory of computation remained a question that both sides regarded as infinitely postponable. Time would tell, the questioners were assured, whether a sow's ear could be turned into a silk purse.[10]

It was inevitable that with computation as the metaphor of the new cognitive science and with computability as the necessary if not sufficient criterion of a workable theory within the new science, the old malaise about mentalism would re-emerge. With mind equated to program, what should the status of mental states be—old-fashioned mental states identifiable not by their programmatic characteristics in a computational system but by their subjective marking? There could be no place for "mind" in such a system—"mind" in the sense of intentional states like believing, desiring, intending, grasping a meaning. The cry soon rose to ban such intentional states from the new science. And surely no book published even in the heyday of early behaviorism could match the antimentalist zeal of Stephen Stich's *From Folk Psychology to Cognitive Science*.[11] There were, to be sure, statesmanlike efforts to make peace between the fuddy-duddy, mentalistic cognitivists and the brave new antimentalists. But they all took the form of either humoring or cajoling the mentalists. Dennett proposed, for example, that we should simply act *as if* people had intentional states that caused them to behave in certain ways; later we'd find out we didn't need such fuzzy notions.[12] Paul Churchland grudgingly admitted that, while it was interestingly problematic why people hung on so tenaciously to their

plainly wrong mentalism, that was something to be explained rather than taken for granted. Perhaps, as Churchland put it, folk psychology seems to describe how things actually go, but how could a belief, desire, or attitude be a *cause* of anything in the physical world—that is, in the world of computation?[13] Mind in the subjective sense was either an epiphenomenon that the computational system outputted under certain conditions, in which case it could not be a cause of anything, or it was just a way that people talked about behavior after it had occurred (also an output), in which case it was just more behavior and simply needed further linguistic analysis. And yes, I must include Jerry Fodor's nativism: it could also be a spinoff of innate processes built into the system, in which case it was an effect rather than a cause.[14]

With the new attack on mental states and intentionality came a related attack on the concept of agency. Cognitive scientists, in the main, have no quarrel with the idea that behavior is directed, even directed toward goals. If direction is governed by the results of computing the utility of alternative outcomes, this is perfectly bearable and, indeed, it is the centerpiece of "rational choice theory." But cognitive science in its new mood, despite all its hospitality toward goal-directed behavior, is still chary of a concept of agency. For "agency" implies the conduct of action under the sway of intentional states. So action based on belief, desire, and moral commitment—unless it is purely stipulative in Dennett's sense—is now regarded as something to be eschewed by right-minded cognitive scientists. It is like free will among the determinists.[15] There were brave holdouts against the new anti-intentionalism, like the philosophers John Searle and Charles Taylor, or the psychologist Kenneth Gergen, or the anthro-

pologist Clifford Geertz, but their views were marginalized by the majoritarians of mainstream computationalism.[16]

I am fully aware that I may be giving an exaggerated picture of what happened to the cognitive revolution once it became subordinated to the ideal of computability in the edifice of cognitive science. I note that whenever a proper cognitive scientist uses the expression "artificial intelligence" (even if it is only once), it is almost invariably followed by the capitalized initials "AI" in parentheses: "(AI)." I take this act of abbreviation to indicate one of two things. The abbreviated form suggests the shortening required by Zipf's Law: the length of a word or expression is inverse to its frequency—"television" eventually reduced to "TV"—with the abbreviation "(AI)" celebrating its comparable ubiquitousness and market penetration. The boast of AI is that it is about *all* mindlike artifacts, even about mind itself, if mind only be considered as yet another artifact, one that conforms to principles of computation. Or the abbreviation, on the other hand, may be a sign of embarrassment: either because there is an aura of obscenity about the artificialization of something so natural as intelligence (in Ireland, by the way, AI is the embarrassed abbreviation for artificial insemination), or because AI is an abbreviation of what, in its full form, might seem an oxymoron (the liveliness of intelligence coupled with the flatness of artificiality). The implied boast of Zipf's Law and the embarrassment of cover-up are both merited. There is no question that cognitive science has made a contribution to our understanding of how information is moved about and processed. Nor can there be much doubt on reflection that it has left largely unexplained and even somewhat obscured the very large issues that inspired the cognitive revolution in the first place. So let us

return to the question of how to construct a mental science around the concept of meaning and the processes by which meanings are created and negotiated within a community.

II

Begin with the concept of culture itself—particularly its constitutive role. What was obvious from the start was perhaps too obvious to be fully appreciated, at least by us psychologists who by habit and by tradition think in rather individualistic terms. The symbolic systems that individuals used in constructing meaning were systems that were already in place, already "there," deeply entrenched in culture and language. They constituted a very special kind of communal tool kit whose tools, once used, made the user a reflection of the community. We psychologists concentrated on how individuals "acquired" these systems, how they made them their own, much as we would ask how organisms in general acquired skilled adaptations to the natural environment. We even became interested (again in an individualistic way) in man's specific innate readiness for language. But with a few exceptions, notably Vygotsky, we did not pursue the impact of language use on the nature of man as a species.[17] We were slow to grasp fully what the emergence of culture meant for human adaptation and for human functioning. It was not just the increased size and power of the human brain, not just bipedalism and its freeing of the hands. These were merely morphological steps in evolution that would not have mattered save for the concurrent emergence of shared symbolic systems, of traditionalized ways of living and working together—in short, of human culture.

The divide in human evolution was crossed when culture