2. CHOMSKY CONTRA DARWIN: FOUR EPISODES

Chomsky, one might think, would have everything to gain by grounding his controversial theory about a language organ in the firm foundation of evolutionary theory, and in some of his writings he has hinted at a connection. But more often he is skeptical.

—STEVEN PINKER 1994, p. 355

In the case of such systems as language or wings it is not easy even to imagine a course of selection that might have given rise to them.

—NOAM CHOMSKY 1988, p. 167

A sizeable gulf in communication still exists between cognitive scientists who entered the field from AI or from the study of problem solving and concept-forming behavior, on the one side, and those who entered from a concern with language, on the other. . . . When the uniqueness of language processes as a human faculty is emphasized, as it has been by Chomsky . . . , the gulf becomes wider.

—HERBERT SIMON and CRAIG KAPLAN 1989, p. 5

On September 11, 1956, at MIT, three papers were presented at a meeting of the Institute for Radio Engineers. One was by Allen Newell and Herbert Simon (1956), "The Logic Theory Machine," and in it they showed, for the first time, how a computer could prove nontrivial theorems of logic. Their "machine" was the father (or grandfather) of their General Problem Solver (Newell and Simon 1963), and the prototype for the computer language Lisp, which is to Artificial Intelligence roughly what the DNA code is to genetics. The Logic Theory Machine is a worthy rival of Art Samuel's checkers program for the honor of AI-Adam. Another paper was by the psychologist George A. Miller, "The Magical Number Seven, Plus or Minus Two," which went on to be one of the classic papers inaugurating the field of cognitive psychology (Miller 1956). The third paper was by a twenty-seven-year-old Junior Fellow at Harvard, Noam Chomsky, "Three Models for the Description of Language" (1956). Retrospective coronations are always a bit arbitrary, as we have seen several times, but Chomsky's talk to the IRE is as good an event as any to mark the birth of modern linguistics. Three major new scientific disciplines born in the same room on a single day—I wonder how many in the audience had the sense that they were participating in a historic event of such proportions. George Miller did, as he tells us in his account (1979) of that meeting. Herbert Simon's own retrospective view of the occasion has changed over the years. In his 1969 book, he drew attention to the remarkable occasion and said of it (p. 47): "Thus the two bodies of theory [linguistics and Artificial Intelligence] have had cordial relations from an early date. And quite rightly, for they rest conceptually on the same view of the human mind." If only that were true! By 1989, he could see how the gulf had widened.

Not many scientists are great scientists, and not many great scientists get to found a whole new field, but there are a few. Charles Darwin is one; Noam Chomsky is yet another. In much the way there was biology before Darwin—natural history and physiology and taxonomy and such—all united by Darwin into what we know as biology today, so there was linguistics before Chomsky. The contemporary scientific field of linguistics, with its subdisciplines of phonology, syntax, semantics, and pragmatics, its warring schools and renegade offshoots (computational linguistics in AI, for instance), its subdisciplines of psycholinguistics and neurolinguistics, grows out of various scholarly traditions going back to pioneer language sleuths and theorists from the Grimm brothers to Ferdinand de Saussure and Roman Jakobson, but it was all unified into a richly interrelated family of scientific investigations by the theoretical advances first proposed by one pioneer, Noam Chomsky. His slender 1957 book, *Syntactic Structures*, was an application to natural languages such as English of the results of an ambitious theoretical investigation he had undertaken into yet another region of Design Space: the logical space of all possible algorithms for generating and recognizing the sentences of all possible languages. Chomsky's work followed closely in the path of Turing's purely logical investigations into the powers of what we now call computers. Chomsky eventually defined an ascending scale of types of grammars or types of languages—the Chomsky Hierarchy, on which all students of computation theory still cut their teeth—and showed how these grammars were interdefinable with an ascending scale of types of automata or computers—from "finite state machines" through "push-down automata" and "linear bounded machines" to "Turing machines."

I can vividly remember the shock wave that rolled through philosophy when Chomsky's work first came to our attention a few years later. In 1960, my sophomore year at Harvard, I asked Professor Quine what critics of his views I should be reading. (I considered myself at the time to be an anti-Quinian of ferocious conviction, and was already

beginning to develop the arguments for my senior thesis, attacking him. Anybody who was arguing against Quine was somebody I had to know about!) He immediately suggested that I should look at the work of Noam Chomsky, an author few in philosophy had heard of at the time, but his fame soon engulfed us all. Philosophers of language were divided in their response to his work. Some loved it, and some hated it. Those of us who loved it were soon up to our eyebrows in transformations, trees, deep structures, and all the other arcana of a new formalism. Many of those who hated it condemned it as dreadful, philistine *scientism*, a clanking assault by technocratic vandals on the beautiful, unanalyzable, unformalizable subtleties of language. This hostile attitude was overpowering in the foreign-language departments of most major universities. Chomsky might be a professor of linguistics at MIT, and linguistics might be categorized, there, as one of the humanities, but Chomsky's work was science, and science was the Enemy—as every card-carrying humanist knows.

> Sweet is the lore which Nature brings;
> Our meddling intellect
> Misshapes the beauteous forms of things:—
> We murder to dissect.

Wordsworth's Romantic view of the scientist as murderer of beauty seemed perfectly embodied by Noam Chomsky, automata theorist and Radio Engineer, but it is a great irony that he was all along the champion of an attitude towards science that might seem to offer salvation to humanists. As we saw in the previous section, Chomsky has argued that science has limits, and, in particular, it stubs its toe on the mind. Discerning the shape of this curious fact has long been difficult, even for those who can handle the technicalities and controversies of contemporary linguistics, but it has long been marveled at. Chomsky's notorious review (1959) slamming B. F. Skinner's *Verbal Behavior* (1957) was one of the founding documents of cognitive science. At the same time, Chomsky has been unwaveringly hostile to Artificial Intelligence, and has been so bold as to entitle one of his major books *Cartesian Linguistics* (1966)—almost as if he thought the anti-materialistic dualism of Descartes was going to come back in style. Whose side was he on, anyway? Not on Darwin's side, in any case. If Darwin-dreaders want a champion who is himself deeply and influentially enmeshed within science itself, they could not do better than Chomsky.

This was certainly slow to dawn on me. In March 1978, I hosted a remarkable debate at Tufts, staged, appropriately, by the Society for Philosophy and Psychology.[7] Nominally a panel discussion on the foundations and prospects of Artificial Intelligence, it turned into a tag-team rhetorical wrestling match between four heavyweight ideologues: Noam Chomsky and Jerry Fodor attacking AI, and Roger Schank and Terry Winograd defending it. Schank was working at the time on programs for natural language comprehension, and the critics focused on his scheme for representing (in a computer) the higgledy-piggledy collection of trivia we all know and somehow rely on when deciphering ordinary speech acts, allusive and truncated as they are. Chomsky and Fodor heaped scorn on this enterprise, but the grounds of their attack gradually shifted in the course of the match, for Schank is no slouch in the bully-baiting department, and he staunchly defended his research project. Their attack began as a straightforward, "first-principles" condemnation of conceptual error—Schank was on one fool's errand or another—but it ended with a striking concession from Chomsky: it just might turn out, as Schank thought, that the human capacity to comprehend conversation (and, more generally, to think) was to be explained in terms of the interaction of hundreds or thousands of jerry-built gizmos, but that would be a shame, for then psychology would prove in the end not to be "interesting." There were only two interesting possibilities, in Chomsky's mind: psychology could turn out to be "like physics"—its regularities explainable as the consequences of a few deep, elegant, inexorable laws—or psychology could turn out to be utterly lacking in laws—in which case the only way to study or expound psychology would be the novelist's way (and he much preferred Jane Austen to Roger Schank, if that were the enterprise).

A vigorous debate ensued among the panelists and audience, capped by an observation from Chomsky's colleague at MIT Marvin Minsky: "I think only a humanities professor at MIT could be so oblivious to the third 'interesting' possibility: psychology could turn out to be like engineering." Minsky had put his finger on it. There is something about the prospect of an engineering approach to the mind that is deeply repugnant to a certain sort of humanist, and it has little or nothing to do with a distaste for materialism or science. Chomsky was himself a scientist, and presumably a materialist (his "Cartesian" linguistics did not go *that* far!), but he would have no truck with engineering. It was somehow beneath the dignity of the mind to be a gadget or a collection of gadgets. Better the mind should turn out to be an impenetrable mystery, an inner sanctum for chaos, than that it should turn out to be the sort of entity that might

yield its secrets to an engineering analysis!

Though I was struck at the time by Minsky's observation about Chomsky, the message didn't sink in. In 1980, Chomsky published "Rules and Representations" as a target article in *Behavioral and Brain Sciences;* and I was among the commentators. The contentious issue, then and now, was Chomsky's insistence that language competence was largely innate, not something that a child could properly be said to *learn*. According to Chomsky, the structure of language is mostly fixed in the form of innately specified rules, and all the child does is set a few rather peripheral "switches" that turn him into an English-speaker instead of a Chinese-speaker. Chomsky says the child is *not* a sort of general-purpose learner—a "General Problem Solver," as Newell and Simon would say—who must figure out what language is and learn to engage in it. Rather, the child is innately equipped to speak and understand a language, and merely has to rule our certain (very limited) possibilities and rule in certain others. That's why it is so effortless, according to Chomsky, for even "slow" children to learn to speak. They aren't really learning at all, any more than birds learn their feathers. Language, and feathers, just *develop* in species ordained to have them, and are off limits to species that lack the innate equipment. A few developmental triggers set the language-acquisition process in motion, and a few environmental conditions subsequently do some minor pruning or shaping, into whichever mother tongue the child encounters.

This claim has encountered enormous resistance, but we can now be sure that the truth lies much closer to Chomsky's end of the table than to that of his opponents (for the details, see the defenses of Chomsky's position in Jackendoff 1993 and Pinker 1994). Why the resistance? In my *BBS* commentary—which I presented as a constructive observation, not an objection—I pointed out that there was one reason to resist this that was perfectly reasonable, even if it was only a reasonable *hope.* Just like the biologists' resistance to "Hoyle's Howler," the hypothesis that life didn't begin on Earth but began somewhere else and migrated here, the psychologists' resistance to Chomsky's challenge had a benign explanation: if Chomsky was right, it would just make the phenomena of language and language acquisition that much harder to investigate. Instead of finding the learning process going on before our eyes in individual children, where we could study it and manipulate it, we would have to "pass the buck to biology" and hope that the biologists could explain how our *species* "learned" to have language competences built in at birth. This was a much less tractable research program. In the case of Hoyle's hypothesis, one could imagine

> arguments that fixed a maximal speed of mutation and selection and showed that there had not been enough time on Earth for the *whole* process to have occurred locally.
>
> Chomsky's arguments, from the poverty of the stimulus and the speed of language acquisition, are analogous; they purport to show that there must have been *large* gifts of design in the infant if we are to explain the speedy development of the mature competence. And while we can take solace in the supposition that we may someday be able to confirm the presence of these innate structures by direct examination of the nervous system (like finding fossils of our extraterrestrial ancestors), we will have to accept the disheartening conclusion that a larger portion than we had hoped of *learning theory*, considered in its most general form as the attempt to explain the transition from utter ignorance to knowledge, is not the province of psychology at all, but rather of evolutionary biology at its most speculative. [Dennett 1980.]

To my surprise, Chomsky missed the point of my commentary. Whereas he himself had offered reflections on what would make psychology "interesting," he couldn't see how there might be something "disheartening" to psychologists in the discovery that they might have to pass the buck to biology. Years later, I finally realized that the reason he didn't see what I was driving at was that although he insisted that the "language organ" was innate, this did *not* mean to him that it was a product of natural selection! Or at least not in such a way as to permit biologists to *pick up* the buck and analyze the way in which the environment of our ancestors had shaped the design of the language organ over the eons. The language organ, Chomsky thought, was *not* an adaptation, but . . . a mystery, or a hopeful monster. It was something that *perhaps* would be illuminated some day by physics, but not by biology.

> It may be that at some remote period a mutation took place that gave rise to the property of discrete infinity, perhaps for reasons that have to do with the biology of cells, to be explained in terms of properties of physical mechanisms, now unknown. . . . Quite possibly other aspects of its evolutionary development again reflect the operation of physical laws applying to a brain of a certain degree of complexity. [1988, p. 170.]

How could this be? Many linguists and biologists have tackled the problems of the evolution of language, using the

same methods that have worked well on other evolutionary puzzles, and getting results, or at least what seem to be results. For instance, at the most empirical end of the spectrum, work by neuroanatomists and psycholinguists has shown that our brains have features lacking in the brains of our closest surviving relatives, features that play crucial roles in language perception and language production. There is a wide diversity of opinion about when in the last six million years or so our lineage acquired these traits, in what order, and why, but these disagreements are as amenable to further research—no better and no worse off—than disagreements about whether the archaeopteryx flew, for instance. On the purely theoretical front, and casting the net much more widely, conditions for the evolution of communication systems in general have been deduced (e.g., Krebs and Dawkins 1984, Zahavi 1987), and the implications are being explored in simulation models and empirical experiments.

We saw in chapter 7 some of the ingenious speculations and models that have been directed at the problem about how life bootstrapped itself into existence, and there is a similar bounty of clever ideas about how language must have got going. There is no question that the origin of language is theoretically a much easier problem than the origin of life; we have such a rich catalogue of not-so-raw materials with which to build an answer. We may never be able to confirm the details, but if so this will not be a mystery but only a bit of irreparable ignorance. Some particularly abstemious scientists may be reluctant to devote time and attention to such far-flung exercises in deductive speculation, but that does not appear to be Chomsky's position. His reservations are directed not to the likelihood of success but to the very point of the enterprise.

> It is perfectly safe to attribute this development [of innate language structures] to "natural selection", so long as we realize that there is no substance to this assertion, that it amounts to nothing more than a belief that there is some naturalistic explanation for these phenomena. [Chomsky 1972, p. 97.]

There have long been signs, then, of Chomsky's agnosticism—or even antagonism—towards Darwinism, but many of us have found them hard to interpret. To some, he appeared to be a "crypto-creationist," but that didn't seem very plausible, especially since he had the endorsement of Stephen Jay Gould. Remember the linguist Jay Keyser's appeal (on page 279) to Gould's term "spandrel" to describe how language came to be? Keyser probably got his terminology from his colleague Chomsky, who got it from Gould, who in return has avidly endorsed Chomsky's view that language didn't really evolve but just rather suddenly arrived, an inexplicable gift, at best a byproduct of the enlargement of the human brain.

> Yes, the brain got big by natural selection. But as a result of this size, and the neural density and connectivity thus imparted, human brains could perform an immense range of functions quite unrelated to the original reasons for increase in bulk. The brain did not get big so that we could read or write or do arithmetic or chart the seasons—yet human culture, as we know it, depends upon skills of this kind. . . . [T]he universals of language are so different from anything else in nature, and so quirky in their structure, that origin as a side consequence of the brain's enhanced capacity, rather than as a simple advance in continuity from ancestral grunts and gestures, seems indicated. (This argument about language is by no means original with me, though I ally myself fully with it; this line of reasoning follows directly as the evolutionary reading for Noam Chomsky's theory of universal grammar.) [Gould 1989b, p. 14.]

Gould stresses that the brain's growth may not have been due initially to selection for language (or even for heightened intelligence) and that human language may not have developed "as a simple advance in continuity from ancestral grunts," but it does not follow from these suppositions (which we may grant him for the sake of argument) that the language organ is not an adaptation. It is, let us grant, an exaptation, but exaptations are adaptations. Let the remarkable growth of the hominid brain be a "spandrel" in whatever sense Gould or Keyser wishes, and *still* the language organ will be as much an adaptation as the bird's wing! No matter how suddenly the punctuation occurred that jogged our ancestors abruptly to the right in Design Space, it was still a gradual design development under the pressure of natural selection—unless it was indeed a miracle or a hopeful monster. In short, although Gould has heralded Chomsky's theory of universal grammar as a bulwark against an adaptationist explanation of language, and Chomsky has in return endorsed Gould's antiadaptationism as an authoritative excuse for rejecting the obvious obligation to pursue an evolutionary explanation of the innate establishment of universal grammar, these two authorities are supporting each other over an abyss.

In December 1989, the MIT psycholinguist Steven Pinker and his graduate student Paul Bloom presented a paper, "Natural Language and Natural Selection," to the Cognitive Science Colloquium at MIT. Their paper, which has itself subsequently appeared as a target article in *Behavioral and Brain Sciences*, laid down the gauntlet:

> Many people have argued that the evolution of the human language faculty cannot be explained by Darwinian natural selection. Chomsky and Gould have suggested that language may have evolved as the by-product of selection for other abilities or as a consequence of as-yet unknown laws of growth and form. . . . [W]e conclude that there is every reason to believe that a specialization for grammar evolved by a conventional neo-Darwinian process. [Pinker and Bloom 1990, p. 707.]

"In one sense," Pinker and Bloom said (p. 708), "our goal is incredibly boring. All we argue is that language is no different from other complex abilities such as echolocation and stereopsis, and that the only way to explain the origin of such abilities is through the theory of natural selection." They arrived at this "incredibly boring" conclusion by a patient evaluation of various analyses of multifarious phenomena that show beyond a reasonable doubt—surprise, surprise—that the "language organ" must indeed have evolved many of its most interesting properties as adaptations, just as any neo-Darwinian would expect. The response from the audience at MIT was anything but boring, however. Chomsky and Gould had been scheduled to reply, so there was a standing-room-only crowd.[8] The level of hostility and ignorance about evolution that was unabashedly expressed by eminent cognitive scientists on that occasion shocked me. (In fact, it was reflecting on that meeting that persuaded me I could no longer put off writing this book.) So far as I know, no transcript of that meeting exists (the commentaries in *BBS* include some of the themes raised at the meeting), but you can recover something of the flavor by contemplating Pinker's list (personal communication) of the ten most amazing objections he and Bloom have fielded since drafts of their paper began to circulate. Versions of most of them, if memory serves me, were expressed at the MIT meeting:

(1) Color vision has no function; we could tell red from green apples using intensity cues.
(2) Language is not designed for communication at all: it's not like a watch, it's like a Rube Goldberg device with a stick in the middle that you can use as a sundial.
(3) Any argument that language is functional could be made with equal plausibility and force when applied to writing in sand.
(4) The structure of the cell is to be explained by physics, not evolution.
(5) Having an eye calls for the same kind of explanation as having mass, because just as the eye lets you see, mass prevents you from floating into space.
(6) Hasn't that stuff about insect wings refuted Darwin?
(7) Language can't be useful; it's led to war.
(8) Natural selection is irrelevant, because we now have chaos theory.
(9) Language couldn't have evolved through selection pressure for communication, because we can ask people how they feel without really wanting to know.
(10) Everyone agrees that natural selection plays some role in the origin of the mind but that it cannot explain every aspect—thus there is nothing more to say.

Are Gould and Chomsky responsible for the bizarre convictions of some of their supporters? This question has no simple answer. More than half of the items on Pinker's list have a clear ancestry in claims that have been made by Gould (numbers 2, 6, and 9 in particular) and Chomsky (numbers 4, 5, and 10 in particular). Those who make these claims (including the others on the list) typically present them on the authority of Gould and Chomsky (see, e.g., Otero 1990). As Pinker and Bloom say (1990, p. 708), "Noam Chomsky, the world's greatest linguist, and Stephen Jay Gould, the world's best-known evolutionary theorist, have repeatedly suggested that language may not be the product of natural selection." Moreover—two important dogs that haven't barked—I have yet to witness either Gould or Chomsky attempting to correct these howlers when they arise in the heat of battle. (As we shall see, this is everybody's weakness; I regret that the siege mentality among sociobiologists has led them to overlook—at any rate, neglect to correct—more than a few cases of egregiously bad reasoning by members of their own team.)

One of Darwin's most enthusiastic supporters was Herbert Spencer, coiner of the phrase "the survival of the fittest" and an important clarifier of some of Darwin's best ideas, but also the father of Social Darwinism, an odious

misapplication of Darwinian thinking in defense of political doctrines that range from callous to heinous.[9] Was Darwin responsible for Spencer's misuse of his views? Opinions differ on this. For my part, I excuse Darwin from the truly heroic task of chastising his champion in public, even though I regret that he wasn't more energetic in pursuing private acts of dissuasion or correction. Both Gould and Chomsky have been vigorous proponents of the view that intellectuals *are* responsible for the applications *and likely misapplications* of their own work, so presumably they are at least embarrassed to find themselves cited as the sources of all this nonsense, for they themselves do not hold these views. (It is perhaps too much to expect their gratitude to me for doing their dirty work for them.)

3. Nice Tries

> *In studying the evolution of mind, we cannot guess to what extent there are physically possible alternatives to, say, transformational generative grammar, for an organism meeting certain other physical conditions characteristic of humans. Conceivably, there are none—or very few—in which case talk about evolution of the language capacity is beside the point.*
>
> —NOAM CHOMSKY 1972, p. 98

> *To make progress in understanding all this, we probably need to begin with simplified (oversimplified?) models and ignore the critics' tirade that the real world is more complex. The real world is* always *more complex, which has the advantage that we shan't run out of work*
>
> —JOHN BALL 1984, p. 159

The problem here is how to get the pendulum to stop swinging back and forth so destructively. Time and again we see the same failure of communication. The truly unfortunate communication gap that Simon and Kaplan speak of (in their quotation at the head of the previous section) is the amplified effect of a relatively simple bit of initial misunderstanding. Recall the difference between reductionists and greedy reductionists (chapter 3, section 5): reductionists think everything in nature can all be explained without skyhooks; greedy reductionists think it can all be explained without cranes. But one theorist's healthy optimism is another theorist's unseemly greed. One side proposes an oversimple crane, at which the other side scoffs—"Philistine reductionists!"—declaring, truthfully, that life is much more complicated than that. "Bunch of crazy skyhook-seekers!" mutters the first side, in defensive overreaction. That is what they would mutter if they had the term—but, then again, if both sides had the terms, they might be able to see what the issues really were, and avoid the miscommunication altogether. That is my hope.

What are Chomsky's actual views? If he doesn't think the language organ is shaped by natural selection, what account does he give of its complexities? The philosopher of biology Peter Godfrey-Smith (1993) has recently focused on the family of views that maintain, in one way or another, that "there is complexity in the organism in virtue of complexity in the environment." Since this was one of Herbert Spencer's pet themes, Godfrey-Smith proposes we call any such view "Spencerian."[10] Spencer was a Darwinian—or you could say that Charles Darwin was a Spencerian. In any event, the modern synthesis is Spencerian to its core, and it is the Spencerism of that orthodoxy that is most often attacked in one way or another by rebels. Manfred Eigen and Jacques Monod are both Spencerian, for instance, in their insistence that it is only through environmental selection that molecular function can be specified (chapter 7, section 2; chapter 8, section 3), whereas Stuart Kauffman's insistence that order emerges *in spite of* (environmental) selection expresses an anti-Spencerian challenge (chapter 8, section 7). Brian Goodwin's denial (1986) that biology is a *historical* science is another example of anti-Spencerism, since it is a denial that historical interactions with earlier environments are the source of the complexities to be found in organisms. Gould and Lewontin's (1979) brief dalliance with "intrinsic" *Baupläne* that account for all but the minor trimmings of organism design is yet another.

Chomsky's suggestion that it is physics, not biology (or engineering), that will account for the structure of the language organ is as pure an anti-Spencerian doctrine as you could find. This explains his misconstrual of my friendly suggestion about passing the buck to biology. I was assuming, as a good Spencerian adaptationist, that "genes are the channel through which the environment speaks," as Godfrey-Smith puts it, whereas Chomsky prefers to think of the genes' getting their message from some intrinsic, ahistorical, nonenvironmental source of organization—"physics," we may call it. Spencerians think that even if there are such timeless "laws of form," they could impose themselves on things only through some selectional process or other.

Evolutionary thinking is just one chapter in the history of Spencerianversus–anti-Spencerian thinking. Adaptationism

is a Spencerian doctrine, and so is Skinner's behaviorism, and so, more generally, is any variety of *empiricism.* Empiricism is the view that we furnish our minds with details that all come from the outside environment, via experience. Adaptationism is the view that the selecting environment gradually shapes the genotypes of organisms, molding them so that the phenotypes they command are some near-optimal fit with the encountered world. Behaviorism is the view that what Skinner (1953, especially pp. 129–41) called "the controlling environment" is what "shapes" the behavior of all organisms. Now we can see that Chomsky's famous attack on Skinner was as much an attack on Skinner's Spencerian view *that* the environment shaped the organism as it was on the limitations of Skinner's model of *how* this shaping took place.

Skinner proclaimed that *one simple iteration* of the fundamental Darwinian process—operant conditioning—could account for all mentality, all learning, not just in pigeons but in human beings. When critics insisted that thinking and learning were much, much more complicated than that, he (and his followers) smelled skyhooks, and wrote off the critics of behaviorism as dualists, mentalists, antiscientific know-nothings. This was a mis-perception; the critics—at least the best of them—were simply insisting that the mind was composed of a lot more cranes than Skinner imagined.

Skinner was a greedy reductionist, trying to explain *all* the design (and design power) in a single stroke. The proper response to him should have been: "Nice try—but it turns out to be much more complicated than you think!" And one should have said it without sarcasm, for Skinner's *was* a nice try. It was a great idea, which inspired (or provoked) a half-century of hardheaded experimentation and model-building from which a great deal was learned. Ironically, it was the repeated failures of *another* brand of greedy reductionism, dubbed "Good Old-Fashioned AI" or "GOFAI" by Haugeland (1985), that really convinced psychologists that the mind was indeed a phenomenon of surpassing architectural complexity—much too complicated for behaviorism to describe. The founding insight of GOFAI was Turing's recognition that a computer could be *indefinitely complicated* but that all computers could be made from simple parts. Whereas Skinner's simple parts had been randomly mated stimulus-response pairings that could then be subjected, over and over again, to the selection pressure of reinforcement from the environment, Turing's simple parts were internal data-structures—different "machine states" that could be composed to respond differentially to indefinitely many different inputs, creating input-output behavior of any imaginable sophistication. Which of these internal states were innately specified and which were to be revised by experience was something left to be investigated. Like Charles Babbage (see note 13 of chapter 8), Turing saw that the behavior of an entity need not be any *simple* function of its *own* history of stimulation, since it could have accrued huge amounts of design over the eons, which would permit it to use its internal complexity to mediate its responses. That abstract opening was eventually filled by GOFAI-modelers with contrivances of dazzling complexity that *still* fell comically short of producing human-style cognition.

Today the reigning orthodoxy in cognitive science is that yesterday's simple models of perception, learning, memory, language production, and language understanding are orders of magnitude too simple, but those simple models were often nice tries, without which we would still be wondering how simple it might, after all, turn out to be. It makes sense to err on the side of greedy reductionism, to try for the simple model before wallowing around in complexities. Mendel's simple genetics was a nice try, and so was the rather more complex "bean-bag genetics" it became in the hands of population geneticists, even though it has often relied on such retrospectively outrageous oversimplifications that Francis Crick was tempted to kick it out of science. Graham Cairns-Smith's clay crystals are a nice try, and Art Samuel's checkers-player was a nice try—much too simple, as we learned, but on the right track.

In the earliest days of the computer, Warren McCulloch and W. H. Pitts (1943) proposed a magnificently simple "logical neuron" from which "neural nets" might be woven, and for a while it looked as if perhaps they had broken the back of the brain problem. Certainly, before they made their modest proposal, neurologists were desperately confused about how to think of the brain's activity. One has only to go back and read their brave flounderings, in the more speculative books of the 1930s and 1940s, to see what a tremendous lift neuroscience got from McCulloch and Pitts.[11] They made possible such pioneers as Donald Hebb (1949) and Frank Rosenblatt (1962), whose "Perceptrons" were, as Minsky and Papert soon pointed out (1969), a nice try, but much too simple. Now, several decades later, another wave of more complicated but still usefully simple nice tries, flying the flag of Connectionism, are exploring portions of Design Space left unexamined by their intellectual ancestors.[12]

The human mind is an amazing crane, and there is a lot of design work that has to have been done to build it, and to keep it working and up-to-date now. That is Darwin's "Spencerian" message. One way or another, the history of environmental encounters over the eons (and during the last ten minutes) has shaped the mind you have right now. Some of the work must have been done by natural selection, and the rest by one or another internal generate-and-test

process of the sort we looked at earlier in the chapter. None of it is magic; none of it involves an internal skyhook. Whatever models we propose of these cranes will surely be too simple in one regard or another, but we are closing in, trying out the simple ideas first. Chomsky has been one of the leading critics of these nice tries, dismissing everyone from B. F. Skinner, through such GOFAI mavens and mavericks as Herbert Simon and Roger Schank, to all the Connectionists, and he has always been right that their ideas have been too simple by far, but he has also exhibited a hostility to the *tactic* of trying for simple models that has unduly raised the temperature of the debates. Suppose, for the sake of argument, we grant that Chomsky could see better than anyone else that the mind, and the language organ which plays such a central role in its superiority over animal minds, are structures of a systematicity and complexity that beggar all models to date. All the more reason, one would think, to search for an evolutionary explanation of these brilliant devices. But although Chomsky uncovered for us the abstract structure of language, the crane that is most responsible for lifting all the other cranes of culture into place, he has vigorously discouraged us from treating it as a crane. No wonder yearners for skyhooks have often taken him as their authority.

He is not the only candidate, however. John Searle is another favorite champion of skyhook-seekers, and he is certainly no Chomskian. We saw in chapter 8 (section 4) that Searle has defended a version of John Locke's Mind-first vision, under the banner of Original Intentionality. According to Searle, automata (computers or robots) don't have real intentionality; at best they have mere *as if* intentionality. Moreover, original or real intentionality cannot be composed of, derived from—or, presumably, descended from—mere *as if* intentionality. This creates a problem for Searle, because, whereas Artificial Intelligence says you are *composed of* automata, Darwinism says you are *descended from* automata. It is hard to deny the former if you admit the latter; how could anything born of automata ever be anything but a much, much fancier automaton? Do we somehow reach escape velocity and leave our automaton heritage behind? Is there some threshold that marks the onset of real intentionality? Chomsky's original hierarchy of ever fancier automata permitted him to draw the line, showing that *the minimal complexity* of an automaton capable of generating the sentences of a human language puts it in a special class—still a class of automata, but at least an advanced class. This was not quite enough for Chomsky. As we have just seen, he planted his feet and said, in effect: "Yes, language makes the difference—but don't try to explain how the language organ got designed. It's a hopeful monster, a gift, nothing that could ever be explained."

An awkward position to maintain: the brain is an automaton, but not one we can reverse-engineer. Is this perhaps a tactical mistake? According to Searle, Chomsky took one step too many before planting his feet. He should have denied that the language organ had a structure that could even be described in automaton terms at all. By lapsing into information-processing talk, talk about rules and representations and algorithmic transformations, Chomsky had given a hostage to the reverse engineers. Perhaps Chomsky's heritage as a Radio Engineer is coming back to haunt him:

> Specifically, the evidence for universal grammar is much more simply accounted for by the following hypothesis: There is, indeed, a language acquisition device [LAD] innate in human brains, and LAD constrains the form of languages that human beings can learn. There is, thus, a hardware level of explanation in terms of the structure of the device, and there is a functional level of explanation, describing which sorts of languages can be acquired by the human infant in the application of this mechanism. No further predictive or explanatory power is added by saying that there is in addition a level of deep unconscious rules of universal grammar, and indeed, I have tried to suggest that that postulation is incoherent anyway. [Searle 1992, pp. 244–45.]

According to Searle, the whole idea of *information processing in the brain*, described abstractly in terms of algorithms that exhibit substrate neutrality, is incoherent. "There are brute, blind neurophysiological processes and there is consciousness, but there is nothing else" (Searle 1992, p. 228).

That is certainly biting the bullet, and biting the same bullet as Chomsky, but in a somewhat different spot: Yes of course, the LAD evolved, and so did consciousness (Searle 1992, pp. 88ff.), but Chomsky is right that there is no hope of a reverse-engineering account of either of them. Chomsky is wrong, however, to grant even the coherence of an automaton-level description of the process, for that opens the door to "strong Artificial Intelligence."

If Chomsky's position on the slippery slope is hard to maintain, Searle's has even more awkward consequences.[13] He grants, as we can see in the passage quoted above, that there is a "functional" story to be told about how the brain does its work in language acquisition. There is also, he grants, a "functional" story to be told about how parts of the brain arrive at depth or distance judgments in vision. *"But there is no mental content whatever at this functional level"* (Searle 1992, p. 234, Searle's emphasis). He then puts to himself the following quite reasonable retort from the cognitive

scientists: "the distinction [between "function" talk and "mental content" talk] does not really make much difference to cognitive science. We continue to say what we have always said and do what we have always done, we simply substitute the word 'functional' for the word 'mental' in these cases." (This is in fact what Chomsky has often said, in reply to such criticisms. See, for instance, 1980.) To answer this retort, Searle (1992, p. 238) is obliged to take a step backwards himself: not only is there no information-processing level of explanation for the brain, he says; there is also really no "functional level" of explanation in biology:

> To put the point bluntly, in addition to its various causal relations, the heart does not have any functions. When we speak of its functions, we are talking about those of its causal relations to which we attach some *normative* importance. . . . In short, the actual facts of intentionality contain normative elements, but where functional explanations are concerned, the only *facts* are brute, blind physical facts and the only norms are in us and exist only from our point of view.

It turns out, then, that function talk in biology, like mere *as if*–intentionality talk, is not really to be taken seriously after all. According to Searle, only artifacts made by genuine, conscious human artificers have *real* functions. Airplane wings are really for flying, but eagles' wings are not. If one biologist says they are adaptations for flying and another says they are merely display racks for decorative feathers, there is no sense in which one biologist is closer to the truth. If, on the other hand, we ask the aeronautical engineers whether the airplane wings they designed are for keeping the plane aloft or for displaying the insignia of the airline, they can tell us a brute fact. So Searle ends up denying William Paley's premise: according to Searle, nature does *not* consist of an unimaginable variety of *functioning* devices, exhibiting design. Only human artifacts have that honor, and only because (as Locke "showed" us) it takes a Mind to make something with a function![14]

Searle insists that human minds have "Original" Intentionality, a property unattainable in principle by any R-and-D process of building better and better algorithms. This is a pure expression of the belief in skyhooks: minds are original and inexplicable sources of design, not results of design. He defends this position more vividly than other philosophers, but he is not alone. The hostility to Artificial Intelligence and its evil twin, Darwinism, lies just beneath the surface of much of the most influential work in recent twentieth-century philosophy, as we shall see in the next chapter.

CHAPTER 13: *When generate-and-test, the basic move in any Darwinian algorithm, moves into the brains of individual organisms, it builds a series of ever more powerful systems, culminating in the deliberate, foresightful generation and testing of hypotheses and theories by human beings. This process creates minds that show no signs of "cognitive closure," thanks to their capacity to generate and comprehend language. Noam Chomsky, who created contemporary linguistics by proving that language was generated by an innate automaton, has nevertheless resisted all evolutionary accounts of how and why the language automaton got designed and installed, and has also resisted all Artificial Intelligence attempts to model language use. Chomsky has stood firm against (reverse) engineering, Banked by Gould on one side and Searle on the other, exemplifying the resistance to the spread of Darwin's dangerous idea, and holding out for the human mind as a skyhook.*

CHAPTER 14: *In chapter 8, I sketched an evolutionary account of the birth of meaning, which will now be expanded and defended against the skeptical challenges of philosophers. A series of thought experiments building on the concepts introduced in earlier chapters shows not just the coherence but the inevitability of an evolutionary theory of meaning.*

---

1. This *bon mot* appeared in the *Tufts Daily*, attributed to Johann Wolfgang von Goethe, but I daresay it is a meme of more recent birth.

2. This is an elaboration of ideas I first presented in Dennett 1975. 1 recently discovered that Konrad Lorenz (1973) described a similar cascade of cranes—in different terms, of course.

3. In fairness to Chomsky, all he says is that free will *might* be a mystery. "I am not urging this conclusion, but merely noting that it is not to be ruled out *a priori"* (Chomsky 1975, p. 157). This mild suggestion has been eagerly inflated by others into a scientifically based demonstration!

4. Two other books in the Library are the most compelling "refutations" of these masterpieces, but of course the Library doesn't contain any refutations, properly so-called, of any of the *true* books on its shelves. These hatchet jobs must be merely *apparent* refutations—an example of a fact that must be true but is systematically useless, since we could never tell which books were which, without the help of, say, God. The existence of this sort of fact will become important in chapter 15.

5. Chomsky has in fact revised his earlier views about the nature of language, making a distinction these days between "E-language" (the external—and you

might say eternal—Platonic object, English, in which so many of the books in the Library of Babel are written) and "I-language" (the internal, intensional, idiolect of an individual), and he denies that E-language is a proper object for scientific study, so he would probably object to the straightforward way I have run this objection (Steven Pinker [personal communication]). But there are more devious ways of running the argument and appealing only to the I-language of individuals. Can Chomsky or anyone else give a good reason for believing that any five-hundred-page book of short sentences meeting the I-language standards of any normal, literate individual would be incomprehensible ("in principle") to that person?

6. Fodor has bitten this bullet: "Nobody has the slightest idea how anything material could be conscious. Nobody even knows what it would be like to have the slightest idea about how anything material could be conscious" (Fodor 1992). In other words, if you so much as *think* you understand the question of consciousness, you're mistaken. Take his word for it—and change the subject, please.

7. This account is drawn, with revisions, from Dennett 1988a.

8. As it turned out, Chomsky was unable to attend, and his place was taken by his (and my) good friend Massimo Piatelli-Palmarini (who almost always agrees with Chomsky, and seldom agrees with me!). Piatelli-Palmarini was the optimal understudy; he had cotaught a seminar on cognition and evolution with Gould at Harvard, and was the author of the article (1989) that had first rendered explicit the Gould-Chomsky position on the nonevolution of language. His article had been a major provocation and target of Pinker and Bloom's essay.

9. Spencer's woolly style was the target of William James' mockery in the epigraph for part II (p. 147). Spencer (1870, p. 396) had offered the following definition: "Evolution is an integration of matter and concomitant dissipation of motion; during which the matter passes from an indefinite, incoherent homogeneity to a definite, coherent heterogeneity; and during which the retained motion undergoes a parallel transformation." The memeology of James' marvelous parody is worth recording. I got the quotation from Garrett Hardin, who informs me that he got it from Sills and Merton (1991, p. 104). They in turn cite James' *Lecture Notes* 1880–1897 as their source, but Hardin has tracked down some further details. P. G. Tait (1880, p. 80) gives credit to a mathematician named Kirkman for his "exquisite translation" of Spencer, of which James' version—presumably borrowed from Tait—is a mutation. Kirkman's (presumably) original version: "Evolution is a change from a nohowish, untalkaboutable all-alikeness, to a somehowish and in-general-talkaboutable, not-all-alikeness, by continuous somethingelsifications and stick-togetherations."

10. It is also one of Herbert Simon's pet themes in *Sciences of the Artificial* (1969), so we might call it Simonian—or Herbertian.

11. One of Warren McCulloch's students, himself a major contributor to these early developments, is Michael Arbib, whose crystal-clear early discussion of these issues (1964) inspired me when I was a graduate student, and whose later work (e.g., 1989) has persistently carved out new territories, still underappreciated by many in the trenches, and on the outskirts, in my opinion.

12. Other nice tries have been the neuroscientists' many models of learning as "Darwinian" evolution in the nervous system, going back to the early work of Ross Ashby (1960) and J. Z. Young (1965), and continuing today in the work of such people as Arbib, Grossberg (1976), Changeux and Danchin (1976), and Calvin (1987)—and Edelman (1987), whose work would be a nicer try if he didn't present it as if it were such a saltation in the wilderness.

13. The remainder of this section draws on my review of Searle's book (Dennett 1993c).

14. Given Searle's position on this, one would predict that he should be utterly opposed to my analysis of the power of adaptationist thinking, as presented in chapter 9. He is. I don't know whether he has expressed this view in print, but in several debates with me (Rutgers, 1986; Buenos Aires, 1989), he has expressed the view that my account is exactly backward: the idea that one can hunt for the "free-floating rationales" of evolutionary selection processes is, in his view, a travesty of Darwinian thinking. One of us has unintentionally refuted himself; the identity of the victim is left as an exercise for the reader.