

Reprinted from

COGNITION

Volume 31, No. 1, February 1989

**Evolution, selection and cognition:
From "learning" to parameter setting in biology and in
the study of language***

MASSIMO PIATTELLI-PALMARINI

Center for Cognitive Science, MIT

pp. 1-44



NORTH-HOLLAND AMSTERDAM

Evolution, selection and cognition: From "learning" to parameter setting in biology and in the study of language*

MASSIMO PIATTELLI-PALMARINI

Center for Cognitive Science, MIT

Received December 1987, final revision accepted October 1988

Abstract

Piattelli-Palmarini, M., 1989. Evolution, selection and cognition: From "learning" to parameter setting in biology and in the study of language. *Cognition*, 31:1-44.

Most biologists and some cognitive scientists have independently reached the conclusion that there is no such thing as learning in the traditional "instructional" sense. This is, admittedly, a somewhat extreme thesis, but I defend it here in the light of data and theories jointly extracted from biology, especially from evolutionary theory and immunology, and from modern generative grammar. I also point out that the general demise of learning is uncontroversial in the biological sciences, while a similar consensus has not yet been reached in psychology and in linguistics at large. Since many arguments presently offered in defense of learning and in defense of "general intelligence" are often based on a distorted picture of human biological evolution, I devote some sections of this paper to a critique of "adaptationism," providing also a sketch of a better evolutionary theory (one based on "exaptation"). Moreover, since certain standard arguments presented today as "knock-down" in psychology, in linguistics and in artificial intelligence are a perfect replica of those once voiced

*This work is part of a research project that I have developed at the Center for Cognitive Science of MIT as a visiting scientist in the academic years 1985-86, 86-87 and 87-88, thanks to a grant from the Alfred P. Sloan Foundation and two successive grants from the Kapor Family Foundation. My sincere gratitude is here expressed to these institutions. The central ideas that have guided the present inquiry have been mostly inspired by the work of Noam Chomsky, Jerry Fodor and Jim Higginbotham and by numerous personal exchanges with each of them. Many a lesson have I learned from (or, perhaps, I should say: many parameters have been fixed for me by) Morris Halle, Kenneth Hale, Jay Keyser, Jacques Mehler and Steven Pinker on the cognitive front, and from Stephen J. Gould and Richard C. Lewontin on the evolutionary front.

Much of this material has been presented and discussed during a graduate course that Steve Gould and I taught at Harvard, in the History of Science Department, during two successive Spring Terms in 1987 and 1988. I am grateful to the department and to Steve for having made this possible. Steve Gould, Morris Halle, Steven Pinker and two anonymous reviewers have offered me precious comments on, and cogent criticism of, a previous draft. I also benefited greatly from Dan Dennett's spirited disagreement with almost everything I say here.

Reprint requests should be sent to Massimo Piattelli-Palmarini, Center for Cognitive Science, MIT, 77 Massachusetts Avenue, Cambridge, MA 02139, U.S.A.

by biologists in favor of instruction and against selection, I capitalize on these errors of the past to draw some lessons for the present and for the future.

1. Introduction: A convergence made explicit

This paper attempts to reconstruct a recent and pervasive change in styles of scientific explanation and, as a consequence of this change, a new ranking of plausibilities for competing theories of acquisition. The morale of this story is that, in sharp contrast to what has happened in the past, a selective theory is by now *always* the most, not the least, plausible candidate for every process of acquisition. Two distinct domains of scientific inquiry, namely modern biology and modern cognitive science, seem to have independently reached this interesting conclusion: Simply put, it amounts to saying that there is no such thing as learning. Admittedly, this is a very counter-intuitive claim, which stands in need of provisos and qualifications. I will provide them, while pleading for our acceptance of the conclusion in the main.

The core fact is that, contrary to received wisdom, there is no *known* process, either in biology or in cognition, that literally amounts to learning in the traditional "instructive" sense, that is, to a transfer of structure *from* the environment *to* the organism. I will be arguing in favour of the even more radical thesis that there *cannot* be any such process, because it is nomologically impossible to obtain any such transfer (or "assimilation," or "interiorization"). With no exception, all the mechanisms of acquisition, of stepwise complexification and of novelty-generation that have been unravelled so far in biology and in the cognitive sciences are due to a process of internal selection. This fact is now universally accepted in biology, but is still controversial in the study of mind and language. Although it may still be argued that there *is*, after all, a process of learning by selection and that the general shift from instructive to selective models is just providing a better concept of *learning* (Dehaene, Changeux, & Nadal, 1987; Toulouse, Dehaene, & Changeux, 1986), I will endeavor to show in some detail why I, for one, see no advantage in the preservation of the term "learning." I agree with those who maintain that we would gain in clarity if the *scientific* use of the term were simply discontinued.

In all fairness, the position I am defending here is far from representing a majority view in linguistics and in cognitive science. In these domains, the majority still firmly believes in learning without quotation marks (Hayes & Broadbent, 1988) and in general intelligence, in perfect agreement with a parallel conception of biological evolution pictured as a stepwise adaptive "problem-solving." Again in fairness, the members of this majority are mostly

right if (and only if) their claim refers to principles and assumptions upheld by biologists before the advent of the deep and far-reaching transformations outlined in the first part of this paper.

I suggest that some important lessons for linguistics and cognitive science *can*, indeed, be drawn from contemporary biology, but that the new principles and the new assumptions came to bury learning by instruction and to replace it with learning by selection, a radically different process. What now replaces learning everywhere in biology has nothing to do with a transfer of structure and everything to do with mechanisms of internal selection and filtering affecting a pre-programmed chain of multiple internal recombinations and internal "switches." As we will see, the basic model of acquisition is best captured by the notion of a hierarchical fixation of internal parameters. Since there are many such parameters and each can be "set" on only one of a small number of admissible values (for many linguistic parameters there seem to be just two such possible values), the final outcome of the process is not pre-determined. Selective fixation, therefore, is not only compatible with individual variability, but it provides the best scientific explanation for variability and diversity.

Modern neurobiology countenances a selective and progressive process of "learning," in the sense that the development of neuronal circuits can unfold as a cascade of interlocked sequences of neuronal outgrowths and selections. These very subtle and powerful mechanisms are active both in pre- and post-natal development, in spite of rather drastic differences in what counts, respectively, as "environmental input." What will be of concern to us here is, obviously, the post-natal selective process, whereby the available and relevant information from the community of speakers operates on structures formed to a large extent before birth, or pre-set to grow through a stepwise process of maturation in which certain crucial stages seem to be open to the information-driven fixation of just one binary choice for the value of each of the central linguistic "parameters." In spite of its combinatorial simplicity and its limitation to just two (or very few) values for each parameter, this selective model is fully capable of providing a satisfactory explanation of complexification and novelty.

It is a general feature of selective models that they can indeed *mimic* the alleged process of learning by instruction, at a purely phenomenological level, by relying entirely on very simple "switch-settings" followed by a variety of combinatorial events. Again, contrary to received wisdom, selectivism turns out to be *more* powerful than instructivism in explaining the origins of specificity, variety and novelty (Changeux, 1985; Changeux, Courège, & Danchin, 1973; Edelman, 1987; Toulouse, Dehaene, & Changeux, 1986).

Partly because of this, but mostly because selective theories persist in com-

ing out *true* almost everywhere in biology, biological thinking has become, ever since Darwin, more and more selective and less and less instructive. There is no place, in the new theories of biology, for an interiorization of external templates, for molding, or for the buildup of richer internal structures dictated by external blueprints. The ascertained existence of very rich innate repertoires, of highly specific dispositions and structures already present in the organism *before* any encounter with the outside world, clashes with the very idea of learning as an instructive process. Development, be it somatic or mental, is essentially the result of this cascade of filters, rather than the cumulative outcome of external instructive "reinforcements" acting on elementary associations and then on associations of associations.

Modern immunology, of all domains of biology, offers a particularly striking case: Instructivism had been for a long time privileged in immunology, but had to yield eventually, from the mid-fifties onwards, to a selectivist explanation. Since the intuitions and the arguments that were once deployed by immunologists in defense of instructivism are very similar to those presently mobilized in artificial intelligence and in some versions of linguistics and cognitive science, a summary analysis of this case study is offered.

The theories and models of modern biology also run against other assumptions that are intimately connected with instructivism: the parsimony of nature, the optimality of design, the adaptive origin of every existing trait, the advantage of thrifty genotypes, the coarse and primitive nature of what is innate. We will see why these assumptions have been either abandoned or drastically reinterpreted in biology, on the basis of data and models coming from molecular genetics and from a new picture of evolution which departs considerably from standard Darwinian adaptationism.

In the first part of the paper I have attempted to draw a sketch of these data and these models as they have emerged in biology proper. In the following sections I have attempted to present a parallel sketch of the data and the styles of explanation offered by a modern scientific study of language.

The data on language acquisition have been privileged because of historical and methodological reasons. The so-called cognitive revolution was in fact sparked some 25 years ago by studies in "transformational grammar" (the updated label is now "generative grammar") (Chomsky, 1957, 1966, 1981, 1986a, 1988). Moreover, no other domain of cognition (yet) offers detailed explanations of selective mechanisms that can match in depth and in generality those employed routinely and as a matter-of-fact by linguists for an awesome variety of natural languages. Many intuitively telling linguistic data militate powerfully against the concept of learning by instruction and all the related concepts (imitation, induction, analogy, problem-solving). In spite of

drastic simplifications, I have carefully preserved the original spirit of these linguistic examples, extracted from a wealth of recent and often quite technical research in generative grammar.

Although language appears to constitute a somewhat autonomous domain in our mental life, that is, in modern parlance a "module" (Fodor, 1983), there are good reasons to believe that at least the *style* of explanation established by generative grammar will find close analogies in other domains of cognition as well. There are domains of cognitive science other than language acquisition where this "inverted" plausibility ranking in favour of selective theories has also been adopted, it seems to me, with great profit. Perceptual and conceptual development in early infancy is the most prominent case, but I could not even attempt to adequately summarize this domain (Carey, 1985; Keil, 1979, 1986; Mehler & Fox, 1985; Smith & Medin, 1981; Spelke, 1985, 1988).

My reconstruction as a whole underlines the convergence between this shift in linguistics and cognitive science and an identical, consistent and presently uncontested shift that has taken place in biology over the last 30 years.

The rise of selective theories in generative grammar and in the version of cognitive science I am here endorsing has been independently motivated. It was not derived from biology, but it is certainly compatible with the explanatory style of modern biology. But we shall witness more than just compatibility: The close resemblance that I have endeavored to make explicit here is, in my opinion at least, both impressive and meaningful. After all, many leading researchers in linguistics and in cognitive science have long considered themselves to be studying the unique biological capacities of our species at an appropriately abstract level. This paper is meant to do justice to their claim.

2. Necessity versus proximity

The first conceptual shift I want to present owes its initial momentum to molecular genetics and evolutionary biology. It is of special importance because it represents the transition from continuism to punctuationism and from strict Darwinian adaptationism to the acknowledgement of gratuitous mechanisms of genetic "hitch-hiking." This new picture, thus, explains why *non*-adaptive causes also lie at the very roots of evolutionary change. Since language and cognition probably represent the most salient and the most novel biological traits of our species, and since their present adaptive value has been constantly used as an *explanation* of their origins, it is now important to show that they may well have arisen from totally extra-adaptive

ing out *true* almost everywhere in biology, biological thinking has become, ever since Darwin, more and more selective and less and less instructive. There is no place, in the new theories of biology, for an interiorization of external templates, for molding, or for the buildup of richer internal structures dictated by external blueprints. The ascertained existence of very rich innate repertoires, of highly specific dispositions and structures already present in the organism *before* any encounter with the outside world, clashes with the very idea of learning as an instructive process. Development, be it somatic or mental, is essentially the result of this cascade of filters, rather than the cumulative outcome of external instructive "reinforcements" acting on elementary associations and then on associations of associations.

Modern immunology, of all domains of biology, offers a particularly striking case: Instructivism had been for a long time privileged in immunology, but had to yield eventually, from the mid-fifties onwards, to a selectivist explanation. Since the intuitions and the arguments that were once deployed by immunologists in defense of instructivism are very similar to those presently mobilized in artificial intelligence and in some versions of linguistics and cognitive science, a summary analysis of this case study is offered.

The theories and models of modern biology also run against other assumptions that are intimately connected with instructivism: the parsimony of nature, the optimality of design, the adaptive origin of every existing trait, the advantage of thrifty genotypes, the coarse and primitive nature of what is innate. We will see why these assumptions have been either abandoned or drastically reinterpreted in biology, on the basis of data and models coming from molecular genetics and from a new picture of evolution which departs considerably from standard Darwinian adaptationism.

In the first part of the paper I have attempted to draw a sketch of these data and these models as they have emerged in biology proper. In the following sections I have attempted to present a parallel sketch of the data and the styles of explanation offered by a modern scientific study of language.

The data on language acquisition have been privileged because of historical and methodological reasons. The so-called cognitive revolution was in fact sparked some 25 years ago by studies in "transformational grammar" (the updated label is now "generative grammar") (Chomsky, 1957, 1966, 1981, 1986a, 1988). Moreover, no other domain of cognition (yet) offers detailed explanations of selective mechanisms that can match in depth and in generality those employed routinely and as a matter-of-fact by linguists for an awesome variety of natural languages. Many intuitively telling linguistic data militate powerfully against the concept of learning by instruction and all the related concepts (imitation, induction, analogy, problem-solving). In spite of

drastic simplifications, I have carefully preserved the original spirit of these linguistic examples, extracted from a wealth of recent and often quite technical research in generative grammar.

Although language appears to constitute a somewhat autonomous domain in our mental life, that is, in modern parlance a "module" (Fodor, 1983), there are good reasons to believe that at least the *style* of explanation established by generative grammar will find close analogies in other domains of cognition as well. There are domains of cognitive science other than language acquisition where this "inverted" plausibility ranking in favour of selective theories has also been adopted, it seems to me, with great profit. Perceptual and conceptual development in early infancy is the most prominent case, but I could not even attempt to adequately summarize this domain (Carey, 1985; Keil, 1979, 1986; Mehler & Fox, 1985; Smith & Medin, 1981; Spelke, 1985, 1988).

My reconstruction as a whole underlines the convergence between this shift in linguistics and cognitive science and an identical, consistent and presently uncontested shift that has taken place in biology over the last 30 years.

The rise of selective theories in generative grammar and in the version of cognitive science I am here endorsing has been independently motivated. It was not derived from biology, but it is certainly compatible with the explanatory style of modern biology. But we shall witness more than just compatibility: The close resemblance that I have endeavored to make explicit here is, in my opinion at least, both impressive and meaningful. After all, many leading researchers in linguistics and in cognitive science have long considered themselves to be studying the unique biological capacities of our species at an appropriately abstract level. This paper is meant to do justice to their claim.

2. Necessity versus proximity

The first conceptual shift I want to present owes its initial momentum to molecular genetics and evolutionary biology. It is of special importance because it represents the transition from continuism to punctuationism and from strict Darwinian adaptationism to the acknowledgement of gratuitous mechanisms of genetic "hitch-hiking." This new picture, thus, explains why *non*-adaptive causes also lie at the very roots of evolutionary change. Since language and cognition probably represent the most salient and the most novel biological traits of our species, and since their present adaptive value has been constantly used as an *explanation* of their origins, it is now important to show that they may well have arisen from totally extra-adaptive

mechanisms. This does not bring us (for the time being) any closer to a detailed reconstruction of the evolutionary origins of language and cognition, but it provides, for the first time, a plausible biological “background” explanation of the way language and cognition *actually* are.

The principles governing language appear, in fact, rather “peculiar” (i.e., non-adaptive and non-optimal) and we will see later why they had been independently ascertained to be largely under-determined by adaptive pressures and by overall functional efficacy. A lot of resistance to this line of inquiry is still due to straightforward Darwinian preconceptions, that is, to the alleged impossibility of reconciling the uniqueness and the extreme specificity of our mental setups with what we “know” about biological evolution “in general” (Dennett, 1987; Lieberman, 1984; Millikan, 1984). On grounds of such “impossibility,” we have been accused of “crypto-creationism,” through what seems to be considered a truly knock-down argument against the version of linguistics and of cognitive science I am here defending. My first point will, then, be to show that the new theory of evolution vindicates the high specificity and the partial disfunctionality of many biological traits, language included, and that, therefore, all accusations of crypto-creationism become totally vacuous.

In hindsight, a turning point in reconsidering what Richard Lewontin (1974) has aptly called “the genetic bases of evolutionary change” was a brief note published in *Nature* in 1971 by the human geneticist Arthur Moutant. He stressed the discrepancy between the slow pace of spontaneous genetic change and the fast rate of biological evolution. The *tempo* and *mode* of speciation appeared to Moutant, as they also did to other biologists, too dramatic and too abrupt to be fully accounted for just in terms of a sequence of random point mutations followed by natural selection (Moutant, 1971). In the lower organisms (prokaryotes), a mechanism was already well known whereby certain viruses were responsible for the “transduction” of genetic material; that is, a process of pick-out-here-and-reinsert-elsewhere. It was also well known that transduction was in some cases randomly governed by the sheer proximity of a gene to the insertion site of the virus and not by the function of that gene. Moutant, in the wake of Sokal and Sneath (1963) and Anderson (1970), pleaded for the case of *generalizing* transduction *across species*. Such an additional whole-gene-smuggling mechanism did not only buy enough time, as well as novelty and expediency, but it also disclosed a remarkably simple, totally myopic strategy of the genetic material vis-à-vis evolution. In particular, Moutant’s bold generalization claimed to be able to cross a major evolutionary frontier: no less than the steep gap between vertebrates and invertebrates.

Since 1971, geneticists have grown accustomed to rather massive DNA

rearrangements. They witness, or even produce at whim, lump insertions of whole blocks of DNA into the genome of a variety of species, jumping with impunity over the frontiers of the animal kingdom. By now, they are fully familiar with instances of spontaneous gene “amplification,” whereby the number of identical copies of the *same* gene within the same cell are counted by the hundreds, thousands, tens of thousands or even hundreds of thousands.

Moutant and other geneticists were entitled to feel uneasy: The canonical process of slow, cumulative “point” mutations followed by a progressive selection of the fittest can only be one factor in the process. In the canonical neo-Darwinian story, which was still very much the received view at the time, a new species must arise from many successive (and “successful”) point mutations working in the same direction, but having been *independently* selected in many distinct genetic loci. The problem is that, often, no great advantage is derived by the organism from a mutation *a* in locus *A*, unless that very organism is *also* and by sheer chance endowed with mutation *b* in locus *B*, mutation *c* in locus *C*, and so on. It can well be the case that, although these mutations are *jointly* very advantageous, one can be lethal in the absence of the other.

In retrospect, Moutant’s hypothesis, that is, the introduction of genetic smuggling across species, had at least the merit of providing, for the first time, enough leeway for an additional evolutionary gamble: the occurrence of major evolutionary changes due to the sheer proximity of one gene to another gene. Our updated knowledge of genetic mechanisms, from gene-amplification to transposons, has diminished, almost nullified, the evolutionary role of genetic reshuffling brought about by viruses. Transduction by a virus is nowadays only one of the events that can explain certain rare and rapid “pulse rates” in evolution. Jumping genes, transposons, chromosomal rearrangements, gene amplification, multiple recognition elements and other natural genetic devilries can do the job as well, or better (McClintock, 1984; Santoro, Mermod, Andrews, & Tjian, 1988). The end of the neo-Darwinian orthodoxy has, by now, been consummated in biology.

One of the core tenets of older Darwinism was that random changes are responsible not only for slight departures from an ideal type, but also for full-blown qualitative *novelty*. Modern biology still wholeheartedly stands by this assumption, having only abandoned the unnecessary, separate assumption of *gradualism*. It is not mandatory any more to assume that, between any two successive life-forms, there must have existed plenty of intermediate forms. As we will see, the theory of “punctuated equilibria” (Eldredge & Gould, 1972; Gould & Eldredge, 1977) and the existence of vast genetic reshufflings relieve us of the hypothesis of gradualism. In evolution, novelty

can also come by jumps. In a nutshell, what the new evolutionary theory is saying is that full-blown evolutionary novelty can also suddenly arise, so to speak, *for no reason*, because novelty caused by sheer proximity between genes is not governed by function and it, therefore, eludes strict adaptationism. Moreover, novelty can well consist in impoverishment, rather than enrichment. Evolutionary novelty is often accompanied, at the biochemical level, by a metabolic *loss* and by functional *specialization*. There can be full-blown evolutionary novelty without any biochemical enrichment (for a non-technical reconstruction of this insight, see Jacob, 1987).

I hardly need to stress how important these ascertained evolutionary discontinuities are for the problem of the origins of language and cognition. We do not "have to" find language "precursors" in the apes, simply because we do not "have to" find intermediate forms in the emergence of *all* biological traits. And, in fact, there are no such language precursors in the apes (Premack, 1986; Seidenberg, 1986; Seidenberg & Petitto, 1979). Stephen J. Gould (1984) never ceases to remind us that many incomplete series in the fossil record are incomplete, not because the intermediate forms have been lost *for us*, but because they simply never existed.

We will see next how, thanks to Jacob's notion of evolutionary "tinkering"¹ and thanks to signal contributions by Stephen Jay Gould, Richard C. Lewontin, Niles Eldredge, Joel Cracraft, Steven M. Stanley, Elisabeth Vrba and others, we presently possess an altogether different, much more subtle theory of evolution. For this new trend, I will adopt the label neo-neo-Darwinism. After all, the prefix "neo" can cover *a lot* of novelty.

3. What is new in the second neo-

The first "neo-" of neo-Darwinism signalled, from the thirties onwards, what is now taught in the standard university textbook: a synthetic theory judiciously encompassing the Darwinian principles, the laws of Mendelian genetics, some elementary facts about DNA and the mathematico-statistical treatment of gene frequencies in model populations (for historical reconstructions, see Jacob, 1970, and Mayr, 1982). But even neo-Darwinism, in spite of half a century of distinguished service and much advertising, is still grossly misrepresented in the naive intuitions of many non-biologists. Being, at bottom, based on subtle statistical mechanisms, evolutionary processes cannot be

grasped by our untrained intuitions; we have to train ourselves by analyzing carefully and *quantitatively* many examples.

We know that our uneducated intuitions concerning even the simplest statistical phenomena are largely defective (Kahneman, Slovic, & Tversky, 1982; Shafir, 1988) and by now we should accordingly know that any naive intuition *whatsoever* concerning evolution is more likely to be wrong than right. All this applies already to neo-Darwinism, that is, to the presently received view of evolution, but naive intuitions help us no better when reconsidered in the light of the new theory, of *neo-neo-Darwinism*, which came to complexify, refine and enrich the received view. Our intuitions have to be, once again, re-educated.

The roles of sheer chance, of sorting and selection are still sovereign, nay, even more so than in the older Darwinian theory, because we must presently countenance a hierarchy of distinct levels,² each governed by its own local rules and quirks, each having some characteristic trend or tendency, each evolving in relative *independence* from all other levels. The notion of fitness and the notion of adaptiveness must be factorized and then relativized to level-specific criteria. The fundamental Darwinian mechanism (a spontaneous generation of diversity followed by selection) is repeated at *all* levels. The point is, however, that diversity is generated at each level owing to mechanisms specific to that level and that the ensuing selective pressures are likewise level-specific. What, ideally, would meet the optimality criteria at one level can generate conflicts with the optimality criteria acting at other levels. The emergent compatibility, across all levels, between all these rules and quirks seems more and more to consist in trade-offs. Jacob's notion of "tinkering" well captures this somewhat sloppy tendency. Evolution by tinkering thrives on transient sub-optimal compromises, sometimes followed by long periods of relative stability which are, in turn, followed by rare sudden reorganizations (sudden when pitted against the scale of geological change, but possibly affecting thousands of successive generations). In modern evolutionary theory all this goes under the name of "punctuated equilibria" (Eldredge & Gould, 1972; Gould & Eldredge, 1977).

It is crucial to emphasize that the events and the trends countenanced by

¹In French *bricolage* (Jacob, 1977); earlier intimations of this concept can be found in Charles Dickens—*Great Expectations*—and in the work of the anthropologist Claude Lévi-Strauss (1962).

²A rough-and-ready census of these levels, each characterizing its own unit of selection, must countenance, going from the less to the more inclusive and from the lower to the higher: single DNA bases, single genes, whole gene-families, horizontally related DNA sequences on different chromosomes of the same cell, the cell itself, the genome of a species as a self-stabilizing and consistent structure, the individual, inclusive groups of genes that are identical by descent and that are carried by different individuals, the group, the actually interbreeding population, the potentially interbreeding population (the species), the niche of actually interacting species, the ecosystem. Each level harbors forces, pressures and optimization strategies that are exclusive to it (Gould & Vrba, 1982; Lewontin, 1970, 1983; Stanley, 1975, 1979; Vrba & Gould, 1986).

this new evolutionary theory are even more "noisy" and capricious than the ones countenanced by classic neo-Darwinism. The reason why this specification is crucial is that in certain quarters, for reasons that are mysterious to me, every critique of the canonical neo-Darwinian theory is invariably and obdurately received as an incitation to espouse some brand of neo-Lamarckism. This is always a very bad move and it would be particularly ill-advised in our case. The picture that I am presenting here is far removed from any Lamarckian assumption or implication. Tinkering and sub-optimality are, as I said, to be supplemented with the idea that evolutionary mechanisms acting at one level are relatively independent from those acting at other levels. The very nature of what drives evolutionary changes bars any possibility of evolution being somehow directed, or guided, or oriented.

The pivotal evolutionary role of sheer proximity effects, what has been dubbed by John Maynard Smith "the hitch-hiking effect of a favourable gene" (Maynard Smith, 1978; Maynard Smith & Haigh, 1974), has led to the more recent distinction between genes and traits that are "selected for" and those that are merely "selected" (Sober, 1985; Sober & Lewontin, 1984). Under appropriate *ceteris paribus* clauses, the mechanism of hitch-hiking assigns a high probability to the event that a gene, if close enough to some other gene actively selected for, will also be selected. Since the positioning of genes along a chromosome often bears no relation whatsoever to their function, selection by sheer proximity amounts to a considerable boost in the randomness and the serendipity of evolution.

These genetic mechanisms based on proximity are not adaptive in the standard sense: They take place at the level of the genome and are driven by the higher probability of joint replication of neighboring strings of DNA. Darwinian and neo-Darwinian adaptationism, on the contrary, focuses on the organism as a whole and calibrates the survival value of the phenotypic traits after they have been hitch-hiked.

Adaptation can, indeed, be a result of hitch-hiking, if the traits expressed by the flanking genes, later on, happen to *acquire* some adaptive value (say, because of a change in the environment or in the mode of life, or because of new behavioral strategies that are made possible by the appearance of these very traits). Adaptation, in this picture, can well be an *indirect*, serendipitous process. We witness a "selection for" taking place *on top of* some previous selection driven by sheer proximity. The marked conceptual difference between these indirect processes and classic Darwinian adaptation is expressed by the new term "exaptation" (Gould & Vrba, 1982; Vrba & Gould, 1986).

It is also important to be reminded that many DNA sequences in our genome and in the genome of other species represent dummy sequences, phenotypically inert material also known as "junk DNA" (or "ignorant

DNA" or "selfish DNA") (Doolittle, 1982; Doolittle & Sapienza, 1980; Dover, 1980; Dover & Flavell, 1982). Whether this "junk" fulfils some indirect function or none at all is presently debated. What is known is that it codes for no protein, that is, it is not expressed into any cellular product. Genetic engineers are, accordingly, quite willing and able to delete it altogether from their final products. I wish to at least mention the existence of junk DNA also because it is a witness for the prosecution of the "parsimony of nature," a misguided assumption to which I shall return in what follows.

From the point of view of classic Darwinian adaptationism, these genetic components ought to be totally, irremediably *invisible*. (It is a truism that, if a gene contributes nothing at all, then it can contribute nothing that is adaptive.) The story can be more complicated, however, because these genetic "dummies" can further increase the evolutionary slack: Mere sorting at one level (the DNA structure of the genome) can be followed by selection at another level (the individual or the group). We witness here just one example of an inter-level "upwards causality": selection at one level based on sorting taking place at some lower level.

Proximity effects, hitch-hiking, gene amplification, viral transduction and other intra-genomic processes can occasionally inaugurate a totally unprecedented evolutionary trend, because they can, by sheer chance, open up new functional possibilities. The process of exaptation and the process called genetic "recruitment" (Kingsolver & Koehl, 1985) can, in the fullness of time, and depending on a fortuitous match between new serendipitous traits and environmental-cum-behavioral modifications, confer on a given organism even greater Darwinian fitness for totally unselected reasons.³ Only a theory equipped with these new, rather crucial, conceptual distinctions is able to describe and explain the meanderings of evolutionary histories. It is clear, however, that a number of classic, straightforward Darwinian principles are still alive in the new theory. If many episodes in these new evolutionary histories, many traits or dispositions in many species, are explained through extra-adaptive (exaptive) mechanisms, no trait can be *counter*-adaptive. If any trait were flatly counter-adaptive, it would quickly disappear, and for quite standard Darwinian reasons. In a sense, therefore, all this suggests a *meta*-Darwinian picture of evolution (certainly not an anti-Darwinian one).

³For instance, one such recruitment, according to Kingsolver and Koehl, has probably turned the insect wings into organs for flight, after they had been selected for as organs of thermal exchange. A recruitment which obviously conferred further adaptive advantages to the wing-carrier and derailed selective pressures onto a different quality space.

4. From parsimony to profligacy

Plenty of recent findings in molecular biology make it hard for us to believe in the economy of nature (Brenner, 1984). I have already alluded to dummy sequences and to gene amplification, responsible for a plethora of identically repeated genes and repeated sequences along our chromosomes and those of other species. Their only reason for being there, perhaps, is just that they replicate fast and well, constituting an instance of what is called "selfish DNA." It is their own "selfish" interest, not the organism's interests, that accounts for their presence. How all this, and much besides (Piattelli-Palmarini, 1985) could be consistent with the notion of the parsimony of nature, I cannot see.

The present opposite catch-word in biology is, significantly, "profligacy." Maxwell Cowan (Cowan, Fawcett, O'Leary, & Stanfield, 1984), Jerry Edelman (Edelman, 1982, 1987; Edelman & Mountcastle, 1978; Edelman & Reeke, 1982), Jean-Pierre Changeux (1980, 1983, 1985), Pasko Rakic and Patricia Goldman (Goldman-Rakic, 1985; Rakic, Bourgeois, Eckenhoff, Zecevic, & Goldman-Rakic, 1986)—among others—are emphasizing the selectivity of neuronal networks, the role of massive cellular death and synaptic degeneracy in the normal build-up of the nervous system. Niels K. Jerne (1974, 1984, 1985), Antonio Coutinho (Coutinho, Forni, Holmberg, Ioars, & Vaz, 1984; Varela, Coutinho, Dupire, & Vaz, 1988) and Pierre-André Cazenave (Jerne, Roland, & Cazenave, 1982) are similarly stressing the profligacy, waste and selectivity in the build-up, the stabilization and the tune-up of the immune system. Innate repertoires, as currently conceived, turn out to be truly immense. We seem to receive "from the hands of nature" (in Hume's expression) much more at birth than we ultimately use in the adult stage. Selective stabilization (Changeux et al., 1973; Changeux & Danchin, 1974, 1976) and the notion of "knowing by unlearning" as originally propounded by Jacques Mehler (1974) converge in turning the developmental problem almost upside-down. The key to the developmental problem is, nowadays, one of getting rid of the superfluous through internal selection, and of generating more complex structures through a selective ontogenetic build-up.

Both in biology and in the cognitive sciences there is a principled discontinuation of the search for a meager stock of all-purpose innate "primitives," out of which everything else ought to be "constructed" through a process of "enrichment," mediated by "association" and "learning." As we will see in a moment, through the exemplary case study of immunology, the present trend is to grant a very rich *innate* repertoire and then look for the mechanisms of internal selection (Edelman & Gall, 1969; Jerne, 1984; Varela, Coutinho, Dupire, & Vaz, 1988).

The key role of the environment is not that of supplying more structure, but rather that of steering a selection. The a priori faith in the economy and the parsimony of nature entailed a concept of learning as a construction, a progressive enrichment in, and of, structures. The present, largely a posteriori, ascertainment of the profligacy of nature suggests we should conceive of "learning" (between quotes) as a process of triggering, filtering, parameter-setting, progressive specialization and selection.⁴ These considerations on parsimony and profligacy have been crucial in the shift from instructivism to selectivism, because *profligacy entails selectivism just as much as parsimony entails instructivism*.

A direct consequence of this state of affairs, one that will be useful when accounting for the evolution of mind and language, is that, contrary to a still widespread assumption, the extreme *specificity* of a trait (in both senses of the word "specificity") is perfectly compatible with its *lack* of survival value. Many highly specific structures and traits of living organisms are, in an important way, *under-determined* by survival criteria *alone*. Many traits can be shown to have acquired survival value after they had been selected for other reasons (see footnote 3), or for no reason (this is the very essence of what we, after Gould, have called "exaptation").

Neo-neo-Darwinism is bringing about a crucial shift in assumptions, assigning to generic external constraints a quite peripheral role. With them, also fades into irrelevancy the notion of "problem-solving." There is no problem presented "objectively" by the environment, to which a species has to "find a solution." Quite often, the species *creates* its own problems; they do not

⁴As rightly stressed by J.-P. Changeux, G.M. Edelman and others, the stepwise selection process is not only compatible with increasing complexification during the ontogenesis of the brain, but it constitutes the best scientific *explanation* of developmental complexity and novelty. This explains why some of these neurobiologists, notably Changeux, consider that what they finally have attained is a satisfactory selective theory of *learning* and see, therefore, no reason to abandon the term learning. A cornerstone of Changeux's theory of learning is his notion of "resonance" between the organism's internal pre-representations and objective external state of affairs. A literal resonance is what allegedly accounts for the internal selection of the right pre-representations (Changeux, 1983, 1985; Changeux, Heidmann, & Patte, 1984; Dehaene et al., 1987; Toulouse et al., 1986). In spite of my admiration for Changeux's remarkable and truly pioneering work over many years and notwithstanding many deep intellectual affinities, I am doubtful that a literal resonance between the internal representations and the external events can explain selective learning. I cannot develop here a detailed critique of Changeux's theory of resonance, but it is perhaps worth pointing out that, at least in the case of language, the objective poverty of the stimulus imposes a resonance involving a variety of silent syntactic elements (see below) and, therefore, a projective reconstruction taking us well beyond the information given. Resonance cannot be a *direct* match with the manifest stimulus, but only with the stimulus *as internally reconstituted*. Because any such resonance is vastly under-determined by the incoming linguistic stimulus, what "resonates" is one internal representation with another internal representation. Resonance is, therefore, "caused" by the input only in the sense that it is triggered by the impinging acoustical pattern. The neo-Cartesian objections to input-naturalism (of which Changeux's is a sophisticated version) are still valid, it seems to me: No "objective" resonance can be the agent of selection, unless it boils down to a releasing mechanism (Piattelli-Palmarini, 1988).

pre-exist "out there" in the environment *before* the species appears (Lewontin, 1982, 1983, 1984a).⁵

Adaptation, simplicity, compensation, equilibration, minimal disturbance and all those universal, parsimony-driven forces of which the natural sciences were once so fond, recede into the background. They are either scaled down, at the physico-chemical level, where they still make a lot of sense, or dismissed altogether. Both in biology and in the cognitive sciences the core problem is one of explaining specifically what there is, rather than one of generating, in the abstract, all that *could* be.

We will see in a moment why in the theory of generative grammar and in cognitive science it is considered at least very implausible that structures external to the organism might *possibly* be "internalized" through a "learning" process and why it is similarly claimed that a process of selection, of triggering and parameter-fixation, acting on a vast, profligate and highly articulated repertoire of innate structures may prove to be the most productive explanatory hypothesis (Chomsky, 1957, 1986a; Lightfoot, 1982, 1988; Pinker, 1984; Wexler & Borer, 1986; Wexler & Culicover, 1980). In linguistics and in cognition at large, this is often received as a *paradox*, something so weird and unacceptable as to be unworthy of consideration. Since an identical attitude has been totally defeated in immunology over the last 30 years, I will briefly present this "lesson from immunology," one that I find particularly compelling.

5. The rise of selectivism in immunology: A case study

Long before this point I should, perhaps, have offered definitions for the terms "selective" and "instructive." The reason I did not is that I do not believe in definitions.⁶ What I am going to offer is, I hope, something better: a case study. The case can be introduced through a very elementary analogy, one that I will borrow from Gerald Edelman (Edelman & Gall, 1969).

Jones needs a new suit; he goes to a department store and chooses the one

⁵Here we witness another interesting convergence between the new evolutionary theory and the cognitive sciences: "Problem-solving" becomes an *uninteresting* metaphor everywhere, as does the metaphor of "trial-and-error" (Piattelli-Palmarini, 1988).

⁶The molecular biologist Antoine Danchin has offered one that is as good as any definition can be: "Instructive theories postulate the existence of a causal agent, exterior to the system, that directs its evolution ... Selective theories, on the opposite, leave to *contingent* interactions the *only* driving force that makes living systems evolve "[italics added]" (Danchin, 1987, p. 189). Noam Chomsky has repeatedly emphasized how casual and unsystematic fragments of linguistic input constitute necessary and *sufficient* data allowing the child to "fix the parameters" of his/her native language (Chomsky, 1981, 1986a, 1986b). My characterization of the theory of generative grammar as selective is justified by such, rather deep, affinities.

that fits best and that he likes best. This is a process of selection. If, on the contrary, Jones had commissioned a tailor for a suit to be cut and sewn *sur mesure*, that would have been an instructive process. Let me quickly add that, if there were stores carrying a repertoire of sizes and cuts large enough to fit *exactly* all *possible* body sizes and satisfy every whim, there would be no need for tailor-made outfits. Moreover, in spite of the dramatic difference in processes and procedures, the final result of selection and that of instruction would be *indistinguishable*.

The plausibility of ready-made repertoires was for a long time considered to be, literally, *unthinkable* in immunology. Considerations of "economy" or parsimony that are quite sensible for a department store were unhesitatingly, obviously, applied to nature. For over half a century, the most distinguished immunologists suggested that all "new" highly specific antibodies had to be generated in the organism as a result of an instructive process, not a selective one. The idea was that the external agent (the antigen) dictated to the organism the shape of the "new," specific antibody that could specifically recognize it.

What is strange in this story is that, already around 1897, one of the early giants of the field, Paul Ehrlich, had originally proposed a *selective* theory of antibody formation, according to which the antigen (say, some molecular configuration present on an infecting virus or a bacterium) encountered in the infected organism some antibodies that were there *already* and which *happened*, by sheer chance, to "recognize" the antigen and bind it specifically.⁷ But the kiss of death for *any* such selective theory came soon enough: As of 1912, Karl Landsteiner and collaborators started gathering experimental evidence to the effect that an organism (say a man, a rabbit or a mouse) can develop very *specific* antibodies against *artificial* substances, compounds that have never been met by any organism anywhere in nature.⁸ It was then considered unthinkable (*unvorstellbar*, in the original scientific German wording of the thirties) that the organism could have, ready-made, highly specific antibodies against artificial substances. Pre-formationism and selectivism were considered not just wrong, but inconceivable. The reason was simple enough: Antibodies against artificial substances were over-specific,⁹

⁷For a later review, see Ehrlich (1900).

⁸For a comprehensive synthesis, see Landsteiner (1936).

⁹In my experience, those who are not fully conversant with immunology often wonder at how "perfect" the fit between antigen and antibody must in fact be. It is also my experience that some skepticism concerning the force of this case is expressed by non-immunologists by virtue of their imagining a certain slack in antigen-antibody recognition and a certain amount of degeneracy in the correspondence between a given antigen and "its" antibodies. The full story would be too long and technical, but suffice it to say that the typical antigen is a complex molecule, harboring many distinct sites of recognition (called epitopes) and that each site indeed

over-fancy and over here. They could not *possibly* be innate. It would have been uneconomical for any organism to store thousands and thousands of *very* specific antibodies against substances that are only to be met in the laboratory, but never in nature. If these antibodies *were* innate, then the immune system, so necessary to survival, would be profligate, wasteful, maladaptive. And *this*, plainly, was "inconceivable" (Alexander, 1932; Breinl & Haurowitz, 1932; Haurowitz, 1952; Mudd, 1932; Pauling, 1940; Topley, 1933).

For over half a century, as a result of Landsteiner's experiments, no immunologist could bring himself to believe in a selective theory of *any* sort. It was only in 1955 that Nils K. Jerne revamped a selective explanation and it took quite a few more years (Burnet, 1959; Lederberg, 1959) and lots of very hard evidence before it was finally accepted by all immunologists (Jerne, 1955; for a reconstruction of the story, see his Nobel lecture, 1985). A strict Darwinian interpretation of the "selection of the fittest" pictured as inevitable the progressive loss of all unnecessary genetic and structural loads. If anything seemed *obviously* useless to the organism, this was the possession of specific antibodies against substances that had never been seen in nature. Since these antibodies could have no *conceivable* survival value, could be of *no* adaptive significance, the equally *obvious* explanation for their demonstrable existence could only be a process of "imprinting" or "learning." The assumptions of parsimony, optimality and adaptiveness led to the inevitable conclusion.

This story is a kind of paradigm case and I have told it in some detail elsewhere (Piattelli-Palmarini, 1986). In the present context I will capitalize on it only to give a clear intuitive grasp of the difference between a selective theory and an instructive theory. Ehrlich, then Jerne and nowadays all immunologists, grant that an organism has the genetically built-in capacity to generate a huge repertoire of many millions of antibody types, an *actual* (not just potential) repertoire so huge that every possible molecular shape, natural

fits very tightly and very specifically one kind of antibody. One source of what might be called "degeneracy" is the multiplicity of distinct sites, each recognized by a distinct family of antibodies (by what is called a specific idio-type) and it is thanks to this "degeneracy" that, for instance, vaccines can be obtained. Fixing our attention to only one such site, recognition is next to perfect by *at least* one line of pre-existing antibodies. Upon re-exposure of the same individual organism to the same epitope, because of further selection, there is *de novo* synthesis of new antibodies that have even greater affinity (this is called "affinity maturation"), although antibodies that have lesser affinity will also be synthesized. The crucial point is that every possible local molecular configuration (every possible epitope) will invariably find in each individual organism, upon first exposure, a pre-existing antibody of high affinity. This first-generation antibody will not cross-react with any other molecule. Well, unless such molecule happens to be, to all intents and purposes, structurally indistinguishable from the epitope in question. The whole process is selective through and through and, indeed, both affinity maturation and degeneracy strengthen the case (Hames & Glover, 1988).

or artificial, is bound to encounter at least one antibody that *by chance* happens to recognize it specifically and bind it selectively. The encounters with the outside world, that is, the contact with the various antigens, selectively shift the demography of the populations of cells which produce the various types of antibodies. Nothing is ever "new" to this system; the repertoire of existing antibodies constitutes a "network," an interactive system of "internal images" of all possible external forms, a repertoire which is "complete" and "closed" (Coutinho et al., 1984; Jerne, 1974; Jerne, Roland, & Cazenave, 1982; Varela et al., 1988). If an antigen, *per absurdum*, were indeed new to the system, the system could do nothing at all with it. A really new antigen would be literally invisible to the immune system: The organism would develop no response to it. This is, typically, a *selective* theory.¹⁰

It is important to add that, in the meantime, the intricate mechanisms of multiple genetic recombination, segment-splicing and multiple joining that factually give rise to the innate repertoire of antibodies and to other, functionally related, repertoires of high specificity (for instance the rich and highly specific repertoire of the T Cell receptors) have been unravelled in the most minute details.¹¹ These data make the pervasiveness of selective mechanisms an undisputed *fait accompli* in present-day immunology.

What, instead, every immunologist in the thirties, forties and early fifties believed was that the external antigen, when coming into intimate contact with the antibody-synthesizing machinery of the organism, *dictated* to it the shape of a *new* antibody. The metaphor was that of an "imprint," a "mould," a "template." The innate component of the immune system could *only* be, in their eyes, an array of shapeless "globulins," of molecular waxes ready to receive the imprint from the outside world; potential receivers of forms, not actual specific structures (Pauling, 1940). This was, typically, an *instructive* theory.

The difference in style ought to be evident. It also ought to be evident how each brand of theory rests on quite different assumptions concerning the economy of nature, the optimality of design, the nature of what is innate. Although the case of immunology is still the most conspicuous in the living

¹⁰The full lesson from the case and the stress on a generally shifting balance from instructive to selective theories are very explicit in Jerne (1967); references to Chomsky's notion of generative power are explicit in Jerne (1985); Chomsky (1988) capitalizes on this case; for a recent re-analysis of the case for selective theories, see also Danchin (1987).

¹¹For the very first—and, at that time, mostly unnoticed—intuition, see Dreyer & Bennett (1965); for the Nobel-prize-winning discovery of the actual genetic mechanisms, see Tonegawa (1983); Hozumi & Tonegawa (1976); for recent reviews, see Marrack & Kappler (1987); Williams (1987); Williams & Barclay (1988); and the comprehensive tutorial papers in Hames & Glover (1988).

memory of molecular biologists, similar transitions from instructivism to selectivism have taken place in other sectors of biology.¹²

This summary sketch of the shift from instructive to selective theories in biology proper ought now to provide a meaningful historical and conceptual background to the analogous shift in the study of language and mind. But, before I proceed to do that, I would like to stress that I am not propounding biological reductionism. My main point here was to explain the reasons why the *a priori plausibility* of theories has changed in biology, why selective mechanisms are by now judged by biologists to be *more* plausible *a priori* than instructive ones *even* when no specific mechanism has yet been unravelled in any detail. Mine is, so to speak, an exercise in mental hygiene, not a plea for reducing cognitive science to immunology or neurobiology. The rise of selectivism in linguistics and in cognitive science was neither historically, nor epistemically, suggested by considerations of our biology and of the channels that our biological makeup provides for our interactions with the environment. The principled reasons for excluding instructive mechanisms in language acquisition are grounded on linguistic data and on linguistic explanations. The high plausibility of selective mechanisms in linguistics and in cognitive science is theory-internal, it was not imported from biology. What we have, therefore, are two independent lines of evidence and two independent styles of explanation. The striking fact is that, nonetheless, they converge as much as they do.

6. Darwinian non-explanations of language and mind

It is one thing to assess the current utility of a biological trait, and quite another to explain the origins of this trait in terms of its current utility. Even if a trait *is* useful and actually enhances the life expectancy of the individuals who possess it, this fact does *not* grant the inference that the trait is there *because* it is useful. In this simple consideration lies most of the difference between evolutionary explanations based on adaptation and those based on exaptation. No exaptationist claims that adaptationism is always wrong, but, rather, that adaptationist reconstructions have been vastly and systematically

¹²I will only mention enzymatic "adaptation" (which turned out to be an all-or-none selective switch), enzyme action (with the demise of the older notion of an "induced fit," now replaced by mechanisms of selective stabilization) and, as we saw, the growth of neuronal networks through the selective stabilization of synapses. The rise of selective theories in many domains of biology has been stressed by Changeux (1983), Jacob (1987), Monod (1970); see also several articles in the collective volume edited by Quagliariello, Bernardi, & Ullmann (1987).

overplayed, and that even when adaptive constraints do apply, they are often *very weak*. Adaptive constraints are typically insufficient to discriminate between real cases and an infinity of alternative, incompatible mechanisms and traits which, although abstractly compatible with the survival of a given species, are demonstrably absent. This applies, and with bells on, to perfectly adaptive linguistic structures that could ideally have been present, but which are, as a matter of fact, wrong *for us*. (I will present some examples in what follows.)

So far, none of the quite specific mechanisms, aptitudes and computational procedures posited by generative linguistics minimally lends itself to an adaptationist explanation. What is worse (from the standpoint of the strict adaptationist) is that, the more we understand them, the less plausible any adaptive explanation appears. These linguistic principles and mechanisms are all, obviously, compatible with our survival, but they are not uniquely (nor, for that, even approximately) "determinable" under the adaptationist constraint *alone*.

An important methodological consideration then ensues: *Only* those who espouse an *instructivist* paradigm are in need of direct adaptive constraints. If, on the contrary, one adheres to the thesis of strong innatism, adaptive constraints become *ipso facto* weak, trivial and theoretically irrelevant.¹³

With a metaphor derived from architecture, Gould and Lewontin have characterized exaptive traits as "spandrels."¹⁴ In the new picture, what we state is that the innate, very specific and highly abstract structures governing language and cognition may also be seen as "spandrels," that is, biological traits that have *become* central to our whole existence, but which may well have originally arisen for some purely architectural or structural reason (perhaps overall brain size, or the sheer duplication of pre-existing modules), or as a by-product of evolutionary pressures driven by other functions (say, an increase in connectivity, or in sub-modularization).

Independently, but, as I said, in significant convergence with the new exaptationist evolutionary theory, generative grammar and cognitive science

¹³Once more, the immune system offers an interesting analogy, because Jerne, Coutinho and Cazenave posit that its *initial* function could have been one of internal recognition, and not one of defense. With a deliberately provocative formula, they plead for "an immunology without the antigen" (Coutinho et al., 1984; Jerne et al., 1982; see also Piattelli-Palmarini, 1985, 1986).

¹⁴"Spandrels—the tapering triangular spaces formed by the intersection of two rounded arches at right angles—are necessary architectural by-products of mounting a dome on rounded arches. ... Such architectural constraints abound, and we find them easy to understand because we do not impose our biological biases upon them. ... In a sense, [these necessary architectural elements] represent an 'adaptation', but the architectural constraint is clearly primary. ... Yet evolutionary biologists, in their tendency to focus exclusively on immediate adaptation to local conditions, do tend to ignore architectural constraints and perform just ... an inversion of explanations" (Gould & Lewontin, 1984, pp. 253–255).

have persuasively stressed that adaptive constraints have no role to play in a scientific approach to language and thought. The structures, mechanisms and modes of functioning of our mental "modules" (Fodor, 1983; for reviews, see Garfield, 1987) derive scant illumination from logical necessity, "general intelligence," overall simplicity, cognitive equilibration, communicative efficacy and survival value. We will see in what follows some examples of the explanatory impotence of these generic adaptationist optimization criteria.¹⁵ It is clear today, in hindsight, why the strong innatist-selectivist could totally disregard these constraints: In his scheme, abstracting away from severe pathological cases or from extreme sensory deprivation, things *cannot go wrong*. He is ready to posit much stronger, more specific and more elaborate internal constraints to grant the over-determination of the outcome.

The study of normal children in a variety of communities (Berwick & Weinberg, 1985; Pinker, 1984; Wanner & Gleitman, 1982), matched with study of the congenitally blind (Landau & Gleitman, 1985), the congenitally deaf (Kegl, 1986, 1987; Klima & Bellugi, 1979; Petitto, 1986, 1987) or even subjects having become blind *and* deaf since 19 months of age (studied by Carol Chomsky and co-workers [1986]—see also the famous and deeply moving case of Helen Keller [1902], who became blind and deaf at age three but who managed, notwithstanding, to become a marvelous writer) show the extraordinary over-determination of the process of language acquisition.

As we will see in what follows through simple examples, what the organism receives as input from the surrounding world is too poor, too erratic and too under-constraining to account for the rich, intricate and powerful setup of mental structures that constitute the linguistic "steady state" of the adult. In spite of the "poverty" of the input, the innate constraints grant that there are *many ways* of getting to the steady state: namely, *all* those actual or potential developmental pathways that are *within* what has been aptly called the "genetic envelope" of the species (Changeux, 1980, 1983, 1985).

A daring assumption is being made by the innatist-selectivist, one that has no "logical necessity," but which turns out to be true nonetheless: *Nothing can be wrong inside the envelope and nothing can be right outside*. Since the envelope, in Changeux's terms, includes *potential* as well as actual pathways, something has to be left "open." Innatism is not to be equated with strict determinism. Fluctuations due to the action of the environment will in fact contingently select different actual structures among all the possible ones. There is ample space here for individual variability and for group diversification. In generative grammar this role is captured by the notion of "parameter-

fixation." The principles of language are innately determined *as* parameterized, that is, open to a wide range of pre-assigned combinatorial options. Within the envelope, parameters are severally set to the values specified by the several actual *or possible* human dialects and languages.

"Learning" a specific language (English, Japanese, Italian etc.) is setting these parameters according to a specific set of choices and acquiring the lexicon of that language (a point to which I return shortly). Innatism and selectivism are, therefore, not only compatible with linguistic variability, but constitute the best scientific explanation of this variability. All real and possible human languages have to remain within the "envelope" assigned by our linguistic makeup, that is, ultimately by our genetic makeup. Every choice inside the envelope is "right," no choice outside can *possibly* be. This core linguistic fact clashes with any belief in "transfers of structure."

The points of contact with biology are here more than simply suggestive: If we are within the envelope, the alleged transfer must boil down to a selection of an *already present* capability, and therefore to no transfer at all; if, on the other hand, a structure lies outside the envelope, it would require no less than a magic trick to make it accessible to the organism.

Under the innatist-selectivist assumption, "transfer of structure" and its equivalents (interiorization, assimilation etc.) become very close to a *nomological* impossibility. But, since transfer of structure is the literal interpretation of the notion of learning in the canonical sense, then these notions are barred from linguistics and psychology as strongly as perpetual motion is barred from physics and the inheritance of acquired characters from genetics. This prompts Chomsky (1986a, 1986b, 1988), Fodor (1975, 1983, 1987) and others, I think rightly, to recommend the discontinuation altogether of the scientific use of the term "learning."

This demise of learning (or, at least, of learning by instruction) stems from two sorts of evidence and from two corresponding lines of argument which I would not want to see conflated. One is biological, while the other is learning-theoretical and each could, in principle at least, be true in the absence of the other. We saw how selective mechanisms acting myopically at different levels can *de facto* explain specificity, novelty, complexification, discontinuity and exaptation in phylogenesis. We also saw that individual ontogenetic development at the biological level does not involve instructive learning. In spite of all this, it would still be conceivable that *cognitive* development be based through and through on instructive mechanisms, if the net outcome of phylogenetic exaptation and selective neuronal growth had been an innate predisposition to generate mental contents by somehow "assimilating" external states of affairs. Here the second line of data and arguments comes to the fore (Piattelli-Palmarini, 1988).

¹⁵For a recent and thorough analysis of the problem of optimization in biology, see the essays collected in Dupré (1987).

As we will see in the case of language, the "poverty of the stimulus" is too severe to make such assimilation even conceivable. What we ascertain is rather a rich *projective* re-construction of highly specific linguistic structures, that is, a process that has powerful *internal* sources and is not based on the assimilation of external templates. This is the innatist-selectivist explanation that linguistic evidence makes most plausible (Chomsky, 1988; Osherson, Stob & Weinstein, 1982, 1986; Wexler & Culicover, 1980). Moreover, there are powerful logical and methodological reasons to dismiss instructive learning *in principle* (Fodor, 1975; see also discussions in Piattelli-Palmarini, 1980). What it is my purpose to underline here, from the combination of these two independent lines of data and arguments, is the conclusion that biology *constrains* acquisition theory in exactly the right way. The specificity and the gratuity of linguistic principles make perfect sense under the hypothesis of an exaptive (i.e., non-adaptive) and discontinuous origin of language. What also makes perfect sense is for language acquisition to be a selective process, since we have just ascertained that, for *every* acquisition process compatible with biological growth, a selective mechanism *must* have (as of now) greater *a priori* plausibility.

7. On the gratuity of natural languages: Language as a spandrel

The versions of linguistics and cognitive science that I am defending here study human nature at an appropriate level of abstraction, one that is just different from the levels explored by neurobiology, genetics and physiology, but not *divorced* from them in nature, scope and force. The specificity of their level of analysis makes them immune from biological "reductionism" and their constant and quite legitimate appeal to a principled deeper continuity with biology also makes them immune from a sterile insularity. Adopting this untendentious naturalistic frame of mind, we now turn to see why adaptiveness, even when assorted with simplicity, parsimony, logical necessity and generic consistency, does not buy enough specificity in the domain of language. Moreover, in the study of mind these additional constraints are not only weak (both singly and in combination), but even less plausible than adaptation, nay, sometimes patently inapplicable (for a rigorous treatment of simplicity in natural language, see Barton, Berwick, & Ristad, 1987; Berwick, 1985; Chomsky, 1955, now in Chomsky, 1985).

The study of language has, in fact, disclosed many instances of specificity and gratuity in the design of all natural human languages, but hardly any instance of traits dictated by generic communicative efficiency, or constraints dictated by the laws of pure logic. Just to give an example, a central principle

of universal grammar (UG)—the grammar underlying the particular grammars of all particular human languages—has been postulated by generative linguistics and called the "projection principle" (Chomsky, 1981). It states that, in every grammatically correct sentence of every natural human language, each "thematic role" (a technical notion which subsumes and generalizes those of agent, object, instrument, event etc.) *has* to be discharged by one *and only one* admissible candidate. This peculiar and idiosyncratic principle, one that has no counterpart in artificial languages or in formal logic, is so powerful *for us* that we cannot help generating in our minds these admissible candidates even when we do not hear—or see—them.¹⁶

Another important principle of UG, a sort of "twin" to this projection principle, states that no particle of language can be added to a sentence without being *ipso facto* interpreted by us.¹⁷ This is called the principle of "full interpretation" (for a thorough discussion, see Chomsky, 1986a). It is perfectly obvious that an abstractly conceivable, but humanly inaccessible, symbolic system of communication could be based on totally different principles. There is nothing "adaptive" or "simple" in these principles.

Generative grammar also specifies in detail how syntactic principles and semantic principles are deeply and finely interconnected, how universal semantic principles are routinely and unconsciously mapped by each of us onto properly "parametrized" syntactic principles. These principles and these relations are innately accessible to us, and the occasional, poor and unsystematic input from the community of other speakers is necessary and sufficient to fix the required parameters. Again, there is no logical necessity here,

¹⁶These are called "empty categories" and there are important differences between different languages in the way these empty categories play their role: What is "empty" in English can be "full" in Italian, Spanish or French, where we have manifest sounds for speech components such as "ne," "se," "ci" etc. (Belletti & Rizzi, 1981; Kayne, 1981; Rizzi, 1982). Empty categories can, also, be totally manifest in the sign languages of the deaf (Kegl, 1986, 1987). (Incidentally, this is a paradigmatic example of how different natural languages differently fix a very crucial parameter.)

¹⁷In artificial languages and in mathematics this principle does not hold. For instance the following are trivially equivalent and well-formed expressions:

a (a) {[a]} 1 × a a/1 a+0 a+b-b

But we cannot adopt the same procedure in English (or in other natural languages):

(i) John likes ice cream

is O.K., but not

(ii) * John likes it ice cream he

we cannot just stipulate that "it" refers to ice cream and "he" to John and that, therefore (i) and (ii) are equivalent. We will have to abide by the principles of UG and write something like the following:

(iii) He, John, likes it: the ice cream.

Aside from nuances of stress and emphasis, (iii) and (i) are *now* equivalent. Nay, in a technical sense, (iii) is close enough to a deeper, implicitly and automatically understood structure of (i), called its "logical form" (for in-depth analyses of this notion, see Chomsky, 1981, 1986a; Higginbotham, 1985, 1988; Lycan, 1984; May, 1985; Schein, 1986).

nor mere conventional expediency. Adaptationism cannot even begin to explain why the natural languages that we can acquire and use possess these central features and not very different ones.

It is interesting to see that this gratuity and this causal necessity are evident not only in the syntax, but also in the semantics of natural languages, a domain that has long been considered to be reducible, on the one side, to the universal laws of propositional logic and, on the other, to the hopelessly episodic, whimsical and utility-governed transactions of everyday life.

From a *scientific* point of view, the most salient facts about meaning are context independent. None of the examples I am going to offer in this paper presumes any other knowledge over and above our knowledge of English. Nor do the examples tap any of our expectations about the world, or convey any "information" that we did not have already. At variance with an older conception that made it the business of semantics to study meaning in general, a new scientific discipline by the same name should rather focus *only* on context-independent conditions on meaning, and on the bases for their acquisition by the child (Higginbotham, 1983, 1985, 1986, 1988).

The interrelations between syntax and semantics lie at the core of what are called "grammaticality judgments" and of their interdigitation with judgments on meaning. Semantics interfaces a lot also with judgment *tout court*, and with reasoning. These negotiations can also be context-independent, have deductive force, but still be distinct from logic (in the usual, disciplinary, normative sense). Some perfectly straightforward semantic inferences elude all formalization by logic *strictu sensu*. We all agree that the ambiguous sentence:

(1) John would be happy with two jobs

can either mean that there are two jobs J_1 and J_2 such that each separately would make John happy, or that there is a situation or event such that John simultaneously has two (unspecified) jobs and that, if it obtained, it would make John happy. The second interpretation, in spite of its being obvious, is impossible to derive through the machinery of standard logic.

Another interesting phenomenon is our capacity to access semantic derivations that are quite obvious to us and yet go well beyond what we manifestly hear, or read. We judge whether the actually spoken/written sentence can have a certain meaning and, if not, then produce an *alternative* sentence that *can* have it. This is well exemplified by sentences such as "I saw two people's pictures," from which we are (interestingly) able to derive an underlying (so-called) "logical form": "I saw pictures such that, for two people x, y , each picture was of x and y ." And immediately realize that it *cannot* be the meaning of the original sentence, but that it *can* be the meaning of *another*, closely related, sentence, namely: "I saw pictures of two people."

The explanation of these very interesting, subtle, though "obvious," semantic facts has to appeal to very elusive properties of plurals, of quantifiers, of reference to events versus reference to sets, of quantification over events, versus quantification over sets (for a thorough analysis, see Schein, 1986). Plainly, nothing here is context-dependent, nor derivable from the laws of pure logic, nor dictated by general communicative constraints, or by criteria of "simplicity." The process of semantic interpretation in natural languages (as persuasively stressed by Higginbotham) is equivalent to a sort of "proof," but a kind of proof that is immediately and effortlessly available to every speaker. Unlike the theorems of mathematics and logic, the deductive steps governing our understanding of language must, therefore, already be available to the child, when he/she begins the process of language acquisition. Semantic interpretation is relevantly like theorem-proving, in that it preserves truth in the derivation from true premisses, but it is also relevantly unlike theorem-proving, in that we do not have to "learn" how to do it, nor work at it painstakingly.

These simple sketches ought to suggest already that survival criteria, the need to communicate and plan concerted action, cannot account for our *specific* linguistic nature. Adaptation cannot even begin to explain any of these phenomena. Not even when it is supplemented with criteria of simplicity or elegance. No human language, for instance, forms questions, or passives, by systematically inverting the order of phonemes, though this procedure would be not only adaptive, but also simple and elegant. It would be so in the abstract, that is, if only we possessed the mental equipment required to "learn" and use forward-backward constructions from our earliest infancy. The fact is that we don't possess such equipment, for reasons that have, presumably, everything to do with evolution, but nothing with adaptation.

Other contingent, that is species-specific, limitations of computational or perceptual capacities seem to be responsible for certain performance failures, for our actual inability to carry out some linguistic task that appears abstractly possible even within the principles of universal grammar. For instance, no intellectual shortage seems to be responsible for our failure to process multiple "center embedded" sentences. Other intelligent creatures might very well manage to communicate in some "naturalized" version of LISP, by using zillions of sentences like the following (or much more elaborate versions):

*the dog the mouse the cat chases sees bites

to mean: The mouse that the cat chases sees the dog that bites the [very] cat that chases the mouse.

Free use of sentences so constructed would allow survival as well as, or perhaps even better than, the use of the sentences we actually produce and understand. "Thoughts" rendered in this way would be very "compact,"