

1

Mind Matters: Chomsky's Dangerous Idea

1 Introduction

If you have ever crossed Harvard Yard, as I have done repeatedly on my way to teach the class on which this book is based, you must have seen visitors gathered around the statue of John Harvard. The statue is known as “The statue of the three lies.”¹ The inscription on the stone supporting the statue reads: “John Harvard, Founder, 1638.” In fact, the model for the statue used by the sculptor (Daniel Chester French) was a student, since there was no portrait of John Harvard available. Equally misleading is the title “founder.” Harvard College was not founded by John Harvard; it was named after him. And the College was founded in 1636, not 1638.

It must have been on one of those occasions when I passed by the statue that I toyed with the idea of calling this book “Chomsky’s Dangerous Idea.” As is the case with the Harvard statue (or, for that matter, with terms like the Holy Roman Empire, which, as Voltaire once shrewdly remarked,² was neither holy, nor Roman, nor an empire), the phrase “Chomsky’s dangerous idea” wouldn’t be accurate, but would nonetheless capture the essence of what I am trying to do in this volume. Although the book will not deal with just one idea, and it won’t always be Chomsky’s, and – I assure you – none of the ideas discussed here are dangerous, it is nevertheless quite correct to say that the central argument developed in the following chapters, shorn of all the qualifications and additions that many researchers have brought forward, goes back to Noam Chomsky. Furthermore, Chomsky’s central claim regarding language remains, in the eyes of many, quite controversial.

What’s this central idea? It will take me a book to characterize it properly, but for now, let me try to state it in simple terms. We, as human beings, come equipped biologically with the ability to develop a language, to make linguistic sense of the noise around us. This ability is both extremely rich in its potential and severely constrained, in ways that cannot be understood by simply looking at properties of

the world around us in which we use language. It requires a *mentalist* stance: a readiness to posit principles of the mind that we, as humans, bring to the task of language acquisition and language use. Since the very beginning of his career,³ Chomsky has urged us to look at the mind in a way parallel to how biologists study the heart, the kidney, or the immune system: as a complex organ of the body, where nature and nurture interact to give rise to some of our most distinctive traits as a species. (For this reason, the approach to language Chomsky advocates is often, and quite appropriately, called *biolinguistics*.⁴)

“Chomsky’s dangerous idea” could also serve as an allusion to Dan Dennett’s popular book entitled *Darwin’s dangerous idea*.⁵ This, too, would be apt, as I happen to think that there are quite a few interesting parallelisms between Chomsky and Darwin. Few scientists have had such a lasting influence on their respective disciplines as these two. Both are extremely controversial figures: They are either reviled or revered (much to their dismay). Both reoriented their own respective fields by articulating a vision that promises to yield genuine scientific understanding and shed light on some of the most difficult questions ever raised. Both showed an impressive ability to revise old ideas, and to turn these ideas into testable scientific hypotheses. Both have advocated theories that, upon reflection, make eminent sense. Indeed, many a scientist would regard Chomsky’s view on language and Darwin’s view on life as virtually the only game in town – the very best we have in these domains of inquiry. And yet, in spite of, or perhaps because of the simplicity of their views, quite a few researchers resist them, and they still require lengthy defenses and expositions.⁶ Finally, like Darwin, the implications of Chomsky’s views for human nature, though implicit in the work that made him famous, only became clear in subsequent publications.⁷

In the end I decided against using “Chomsky’s Dangerous Idea” as the title for this book. I wanted something less personal, and more immediately transparent, something with key terms like “language,” “cognition,” “mental,” and “structures,” to serve as *leitmotifs*. These will be the real characters in my story. My overall aim is to convince you that the study of human language offers a unique vantage point from which to understand human nature (“languages are the best mirrors to the mind,” as Leibniz once asserted⁸), from which to do cognitive science, from which to dig deep into what William James called “the science of mental life.”⁹ Linguistics, as practiced along Chomskyan lines, is an essential player in the elaboration of a genuine “Science of Man.”¹⁰

I would be happy if by the end of the book I have managed to convey to the reader some feeling for what is without a doubt one of the most important and exciting paths of discovery that our species has embarked upon: the search for the underlying principles that both govern and constrain what our brain does (i.e., what our mind is). It is an avenue of research that many have explored, but where progress has been very slow. Cognitive science as we know it today is a relatively young science. It is only in the past 50 years that substantial advances have been made, and much – so much! – remains to be discovered. To the novice, this can be both a source of frustration, and excitement. Frustration because even some

of the most basic questions are still up for grabs, leading to the impression that the foundations are not firm, that too much is changing too rapidly, that in the end very little is known, and that much has to be unlearned on a frequent basis. (These are impressions that I have invariably encountered in my teaching, and that I will try to dispel. I'm introducing them now to urge the reader to recognize them and resist them.)

Excitement because even some of the most basic questions are still up for grabs, leading to the impression that fundamental discovery could come from one's own research, that the missing link is within reach, and that this is history very much in the making. (These are impressions that I have invariably encountered in my teaching too, and that I will do my best to cultivate throughout our journey. May the reader never abandon them in the face of difficulty!)

2 I-language

Using language to study the human mind is a very old practice. Since at least Descartes¹¹ in the modern era (but no doubt, the practice goes back much further), it has been standard to treat language as providing privileged access to some of the deepest secrets of the mind. Unfortunately, the view of language that is required to probe cognition is quite different from our common-sense notion of language, and the way we experience it on a daily basis. We tend to think of language as something that is hard to acquire, that varies from place to place, that is inextricably linked to social norms and culture, and so on. And yet I will argue, following Chomsky, that if we are to make progress in the domain of cognition, language must be understood as something that is acquired effortlessly, that is shared by the entire human species and fundamentally the same across cultures, and that is radically dissociated from social norms. This, I realize, will be hard to swallow at first, and part of this chapter will be devoted to rendering this view somewhat plausible. In an attempt to make this change of perspective conspicuous, Chomsky has suggested¹² we distinguish between language as seen from a social/cultural perspective (what he calls "*E(xternal)-language*"), and language as seen from a cognitive perspective (what he calls "*I(nternal)-language*"). Unless otherwise indicated, whenever I use language, I mean "I-language." (This is what the word "in" in the title of the book, *Language in Cognition*, is meant to emphasize.)

I ask the reader to keep the distinction between I-language and E-language firmly in mind because invariably, when I find myself in a casual setting and I mention the fact that I study linguistics, I am asked how many languages I speak. To a linguist like myself, this question is distinctly odd, for many of us (myself included) think there is only one human language on the planet. This is not to deny that there exist many linguistic variants (languages, dialects, etc.), but these are not the primary objects of inquiry to the linguist/cognitive scientist. Asking a linguist how many languages she speaks/knows is a bit like asking how many

numbers a mathematician knows, or how many species a biologist knows, or how many molecules a chemist knows.

A linguist, to the extent that she is interested in finding out how some aspects of the mind work, tries to elucidate the mental capacities in virtue of which humans are able to produce and understand utterances in *any* (human) language variant. Linguistics, in this sense, is like trying to discover the laws of (some aspects of) the mind. Like any other scientist, the linguist will have to learn to go beyond the multitude of languages to get at the fundamental principles that make linguistic experience possible. From this perspective, specific human languages (English, Navajo, Japanese, . . .) are used as tools, convenient entry points to study something more fundamental, the same way biologists study rats, fruitflies, and worms. Here the linguist has the distinct advantage that languages are literally all over. We are surrounded by language(s). For this reason, linguistic facts often appear to be “cheap.” There is no need to deal with dangerous equipment, flammable substances, etc., to get something to work on. But linguistics is not blind data collection. It is an empirical science. One must be as careful with linguistic data collection and interpretation as in a science lab. And herein lies the rub. Many of the data that linguists deal with pertain to aspects of language that we use instinctively, and that we therefore are unaware of. Grammar classes in high school don’t dwell on them, and they are not discussed in popular newspaper columns about language. And because our mastery of these facts is so effortless, we tend to think that the explanation for these facts must be straightforward and intuitive. But as George Gershwin put it, it ain’t necessarily so.

3 A Few Illustrations

Consider the fact that you can readily recognize ambiguous sentences like these (often, contextual information, or special intonation favors one reading, but given enough time all native speakers are able to detect the ambiguity):

- (1) Flying planes can be dangerous.¹³ (either planes that fly are dangerous, or flying these planes is dangerous)

or

- (2) Mary hit the man with the umbrella.¹⁴ (either the man had the umbrella, or Mary used the umbrella to hit the man)

or

- (3) John can ride a bike. (either John is allowed to, or he has the ability to, ride a bike)

Data like these are so mundane that the richness they contain hardly gets noticed. For instance, did you ever pause to marvel at the fact that a sentence like (4) is only two-way ambiguous, and not four-way ambiguous?

(4) John can ride a bike, and Mary can ride a bike, too.

On the face of it, you might have expected that putting two two-way ambiguous sentences together (*John can ride a bike; Mary can ride a bike*) would give rise to a four-way ambiguity, but somehow, it doesn't. Something forces you to understand *Mary can ride a bike* in the same way you understand *John can ride a bike* when the two sentences are combined into one. And no context, no matter how contrived, will make this fact go away. You must interpret members of a conjunction in parallel. Let's call this the Parallelism Constraint.

If you are like my students, you may be tempted to say that the Parallelism Constraint is due to the presence of the words *and* and *too*. But how about (5)?

(5) John can ride a bike, but Mary cannot ride a bike.

Here too, you have to understand the meaning of the word *can* in the same way in both sentences. But it's not clear what word to blame this constraint on this time. I'm sure you have never thought about this fact before (unless you met a linguist at a cocktail party, and he or she tried to impress!), and you were never told how to interpret ambiguous sentences. Facts like those just discussed constitute *tacit* (unconscious) knowledge: knowledge that the great linguist Morris Halle aptly called "unlearned and untaught."¹⁵ It's just part of what it means to be a native speaker of a language. (Here and so often elsewhere in the book, I will use English examples to illustrate the points I want to make, but this is just a matter of convenience. Similar examples could be found in any other language of the world.)

The extent of your tacit knowledge of language does not stop here. We have, in fact, barely scratched the surface. But, if you think that the ability to understand ambiguous sentences is a purely linguistic ability, think again. You can, for example, readily construe the cube to the left of Figure 1.1 from two perspectives: (a) and (b).¹⁶




Figure 1.1 Two visual interpretations of the same outline cube

And here too, juxtaposing two ambiguous objects appears to demand a parallel interpretation. You must see the cube on the left in Figure 1.2 from the same perspective you see the cube on the right.¹⁷ In the same vein, you may have felt in cases like (4) and (5) that the sentences contained redundant material, and you may have felt to desire to shorten them like this: *John can ride a bike, and Mary can, too; John can ride a bike, but Mary cannot.*



Figure 1.2 Viewing objects from the same perspective

This ability to understand material that is not explicitly present is also part of your tacit linguistic knowledge, but it shows up elsewhere in cognition. For example, everyone I have met interprets the visual object in Figure 1.3a as a white strip covering a black strip.¹⁸ You would, I'm sure, be surprised, if there was nothing underneath the white strip; that is, if you found what is in Figure 1.3b. You would be equally surprised if there was something as weird as what is in Figure 1.3c




Figure 1.3 What's underneath the white strip?

underneath the white strip. This shows that your visual system imposes severe constraints on object construal. This is true of your linguistic system too. You would be surprised if I told you that *John can ride a bike but Mary cannot* means something like “John can ride a bike, but Mary cannot drive a car.”

4 A Few More Notions

Linguists and other cognitive scientists are trying to figure out what these constraints on interpretation are, how exactly they work, and why they should be part of our mind. In so doing, linguists practice what is known as *descriptive grammar*, as opposed to *prescriptive grammar*.¹⁹ Prescriptive grammar is what people tend to associate with the term “grammar.” It refers to a normative practice that seeks to dictate proper linguistic usage. I guess you can say that this is the sort of thing that everyone ignores, unless they have to draft an official letter, make a public address, and so on. Prescriptive grammarians will tell you never to end sentences with a preposition, and you’d be right to wonder what those guys are talking about [sic]. Descriptive grammarians do not impose artificial rules that no one follows; they instead focus on normal usage, and try to figure out the underlying capacity that is responsible for what is said, and how it is interpreted. In so doing, descriptive grammarians also focus on what is sometimes called “negative” knowledge, or knowledge of negative facts: things that speakers never do. We already saw a negative fact in the context of the Parallelism Constraint (you *cannot* interpret sentences like *John can ride a bike and Mary can too* as four-way ambiguous). Here is another. Whereas native speakers of English commonly end their sentences with prepositions (*What did you say that Mary was talking about?*), they never say things like *What did you say about that Mary was talking?* For some reason, it is OK (cognitively/descriptively, but not prescriptively) to leave a preposition dangling at the end of a sentence, but it is not OK (cognitively/descriptively, and – irrelevantly – prescriptively) to leave a preposition dangling in the middle of a sentence. Prescriptive grammarians never worry about such negative facts (since no one ever does it, there is nothing to prohibit), but descriptive grammarians use them all the time to discover the limits of linguistic knowledge, the boundaries of our language faculty.

In some cases the limits will not be strictly linguistic, they will be shared with other cognitive abilities. This is arguably the case for the Parallelism Constraint. In such situations, the constraint will be said to be part of our language faculty broadly construed (*Faculty of Language Broad*, or *FLB*). But in some other cases, the negative knowledge will be strictly linguistic (part of our language faculty narrowly construed; *Faculty of Language Narrow*, or *FLN*).²⁰ (Teasing apart the contributions of FLN and FLB is one of the most exciting ways of figuring out how the mind works. For this reason I will devote a whole chapter (Chapter 12) to the distinction.)

One of the clearest examples of such an FLN-constraint comes from the way we interpret pronouns (elements like *he*, *she*, *it*, etc.). Native speakers of English have very sharp intuitions about sentences like (6–9).

- (6) John likes him.
- (7) He likes John.
- (8) John said that he likes Mary.
- (9) He said that John likes Mary.

In a normal context (with normal intonation), *John* and *him* cannot be understood as referring to the same person (meaning “John likes himself”). The same is true of *he* and *John* in (7). But in (8), *John* and *he* can “co-refer,” though not in the minimally different sentence in (9).

Sentences like (10) and (11) (where co-reference between *John* and *he* is possible in both sentences) make clear that whatever the principle at play is, it cannot be something very simple and intuitive. And yet, speakers’ intuitions are remarkably uniform, and consistent.

- (10) When John came in, he sat down.
- (11) When he came in, John sat down.

When analyzing sentences like (6–11), linguists invariably appeal to technical terms like dependent clause, embedded clause, and so on – constructs that I do not want to go into at this stage (I will come back to them in later chapters).²¹ The take-home message for now is that, as far as one can tell, there is no way to account for the paradigm in (6–11) without appealing to strictly grammatical notions – which speakers are not aware of, but which, at some level, must exist if we are to account for how they understand sentences. (Even something like “pronoun” is a concept or category that speakers must rely on to formulate the rule.)

5 Problems and Mysteries

From the discussion so far, it is clear that we are such thoroughly linguistic animals that we hardly realize what a complicated business language is. I readily confess that some of the sentences that linguists focus on are not very frequently uttered, but in so doing, linguists are no different from other scientists, who try to isolate hidden principles by performing artificial, and highly contrived experiments. Linguists are lucky that they can run interesting experiments very quickly, by probing speakers’

intuitions about sentences they may have never heard before. For here lies one of the most amazing facts about our linguistic knowledge: we are able to produce and understand sentences that we have never heard before. This is sometimes referred to as *the creative aspect of language use*.²² The most conspicuous manifestation of it is our ability to produce (and understand) ever-longer sentences:

- (12) John did it.
 Mary said that John did it.
 Bill claimed that Mary said that John did it.
 Sue believes that Bill claimed that Mary said that John did it.
 Harry thinks that Sue believed that Bill claimed that Mary said that John did it.
 ...

And speakers can readily judge sentences as acceptable, unacceptable, good, awkward, impossible, etc. no matter how unfamiliar the sentences are. As Chomsky famously noted, native speakers of English can judge a sentence like *Colorless green ideas sleep furiously* as a grammatically possible sentence of English (though it's hard to know what it could possibly mean), and can readily distinguish it from *Green sleep ideas furiously colorless*, which is not only meaningless, but also utterly impossible (linguists like to call such examples "word salad").

It is such intuitions (technically known as *acceptability judgments*) that linguists rely on in their attempt to uncover the principles responsible for mental structures. I think it is fair to say that few would have suspected the richness that 50 years of intensive research on mental structure have revealed. And I hope to reveal some of that richness here, and, hopefully, convince you that you too can contribute to this science.

Before concluding this chapter with an overview of the rest of the volume, I want to make a cautionary note. Although I hope that the present book will convince you of how much can be learned about human nature by studying human language, it is also important to remain aware of the fact that discovering mental faculties, such as the faculty of language, goes hand in hand with discovering their limits. Although the research program carried out by modern linguists has proven remarkably successful in some respects, it has also proven quite limited in other respects. I do not want to make false promises: This book will provide little, if any, insight into the nature of Consciousness, the Self, Meaning, and other topics that the layman tends to associate with "the Mind."²³

Cognitive science, like science in general, does not intend to be a theory of everything. The path of science is usually one of humility, a journey during which one learns to lower one's expectations, and appreciate how deep one can go into a few simple topics. The great physicist Richard Feynman was right when he said that the reason physics has been successful is because it focuses on a few simple things (leaving the rest to the chemists, who in turn pass it on to the biologists, who in turn have passed it on to the cognitive scientists, who in turn pass it on to

the novelists).²⁴ The Ancient Greeks thought that it would be easier to discover principles of the mind than principles of the heavens, because the mind was somehow “closer” and therefore more readily accessible to us. History has taught us that just the opposite is true. It may well be that, as Descartes suspects, we “do not have intelligence enough”²⁵ to understand some of the things that we are most curious about. The tale that follows emphasizes the existence and richness of cognitive abilities. In so doing, it also highlights their limitations. It may very well be that the very same cognitive abilities that enable us to understand some aspects of the world set us on a wrong path when it comes to other aspects of the world. As Noam Chomsky often remarks,²⁶ the very same organ (the brain) that defines what can be a problem for us (something which we can investigate with some hope of understanding) also defines what must remain a mystery (something which we cannot even turn into a problem). Just like it’s physically impossible for us to do certain things, there are things that are cognitively impossible for us to achieve; that means that there are some questions we will never answer.

6 Organization

With all these preliminary remarks in mind, let me now touch on the structure of the book. I have decided to organize the material around the five questions that have defined linguistic inquiry over the past 50 years. These are:²⁷

- 1 What is the best characterization of our Knowledge of Language?
- 2 How is that Knowledge acquired?
- 3 How is that Knowledge put to use?
- 4 How is that Knowledge implemented in the brain?
- 5 How did that Knowledge emerge in the species?

These five questions touch on many disciplines, beyond traditional areas of linguistics: evolutionary biology, neuroscience, developmental psychology, and more. Linguistic inquiry, guided by these questions, is interdisciplinary science in the strongest sense. It forces scientists to join forces, and to learn from each other.

Inquiry guided by the five questions just listed has made substantial progress. The questions have led to the best theory developed about language and its place in the human mind, and, as we will see in subsequent chapters, they can be taken as a model to study other cognitive abilities. Cognitive science is still very much in its infancy, and many cognitive abilities remain untouched. Recent work on music by Ray Jackendoff and Fred Lerdahl,²⁸ and on morality by John Mikhail, Sue Dwyer, and Marc Hauser,²⁹ have shown that adopting the methodology of linguistics (“pushing the linguistic analogy,” as it is sometimes described) can lead to important insights into the structure of other mental faculties.

At the same time, the answers provided by linguists have reached a certain depth that make them suitable for use as bridges to close the gap between our current understanding of the mind and of the brain. Although Spinoza in the seventeenth century already stated that the mind is the idea of the brain,³⁰ a formulation echoed by Steve Pinker's definition of the mind as what the brain does,³¹ we are in need of linking hypotheses connecting brain and mind. But as researchers like David Poeppel have stressed for years,³² such linking hypotheses are hard to come by because the concepts used in mind science and in brain science are not commensurable. It is hard for neurobiologists to know what to look for in the brain when what they are given by psychologists doesn't quite match what they already understand about the brain. One can only hope that as the tools and results of cognitive science are refined, they will become usable to formulate precise questions about the brain. There are signs, from very recent studies, which I will review in due time, that linguistics may be getting closer to defining an actual research program for neuroscience. It is still very preliminary, but I think there are grounds for optimism.

In sum, linguistics will play three roles in the following pages: it will be used as a theory (of a particular aspect of human cognition, the language faculty), as a model (to investigate other aspects of human cognition), and as a program (to formulate questions about how exactly the brain produces the mind). If I am correct about this, linguistics ought to be seen as one of the core areas of the Science of Man, an area that no serious cognitive scientist can afford to ignore.

To help the reader better appreciate where we are now, and what we can reasonably expect in the future, it is very useful to examine the context of emergence of modern cognitive science. For this reason the next chapter discusses the major factors that led to what is often called the "cognitive revolution"³³ (or "the mind's new science")³⁴ that took place in the 1950s. I then turn to a discussion of phenomena that taken as a whole provide some of the strongest reasons to investigate the biological foundations of language, which motivate modern linguistics.

The next two chapters thus alternate between the very general bird's-eye-view on cognitive science, and the more narrow, sometimes even microscopic, focus on the nature of the language faculty. In so doing I seek to establish a rhythm that I try hard to maintain in subsequent parts of the book, alternating between the laserbeam (focus on the details) and the floodlight (highlight the big picture). Although language will often be foregrounded, the reader should never lose sight of the fact that this is all in aid of a much bigger research project, one that promises to deliver insights into our most fundamental nature as a species. It's the bigger picture that Einstein urged us never to forget while doing our equations.