

Chomsky's naturalism: its scope and limits

Pierre Jacob
Institut Jean Nicod,
UMR 8129,
EHESS/Département d'études cognitives-ENS/CNRS
Paris, France
jacob@ehess.fr

Fifty years ago, Noam Chomsky laid the foundations for a new scientific approach to the human language faculty (HLF), which he called “generative grammar.” Furthermore, his argument that behaviorist explanations of human verbal behavior are inadequate was a major input to the cognitive revolution that took place in the 1960's and that gave rise to the cognitive sciences.¹ Today few contemporary analytic philosophers of mind or language would deny, I think, that Chomsky's work has deeply changed our scientific understanding of human language.

Many philosophers, however, have challenged one or another aspect of Chomsky's framework for investigating the human language faculty. Not only has Chomsky systematically responded to his critics, but he has also produced his own evaluations of their contributions to the understanding of the human mind and human language. As a result and as

¹ Interestingly, in his (2003) paper, the psychologist George Miller, who was also a major actor of the so-called “cognitive revolution,” argues that the revolution in question should rather be seen as a counter-revolutionary response to the behaviorist revolutionary denial of internal psychological intermediaries between sensory inputs and behavioral output.

two recent publications demonstrate,² two gaps now separate Chomsky from the community of analytic philosophers: on the one hand, many philosophers cannot subscribe to what Chomsky calls “methodological naturalism.” On the other hand, unlike some of the philosophers who do subscribe to methodological naturalism, Chomsky rejects what he calls “metaphysical naturalism.”

My goal here is to explore the nature of these two divides. Over the years, Chomsky has come to make an important distinction between two versions of a naturalistic approach to human mind and language: *methodological* naturalism — which he accepts — and *metaphysical* naturalism — which he rejects. In the first section, I will succinctly characterize the conceptual framework created by Chomsky for the scientific study of the HLF. As I will argue in the second section, many (if not all) of the criticisms directed by philosophers at Chomsky’s scientific framework show that, whether intentionally or not, they embrace some version of methodological dualism, which is inconsistent with the methodological naturalist position advocated by Chomsky. But, in my view, the most unexpected, and the most interesting, divide is that which separates Chomsky from the program of some of the philosophers who subscribe to metaphysical naturalism, and whose aim is to “naturalize intentionality.” In the third section, I will examine the question of what prevents Chomsky from accepting metaphysical naturalism. In the fourth section, I will examine Chomsky’s reservations about the program of naturalizing intentionality.

1. The scope and limits of scientific investigation of the language faculty

“Naturalistic” investigation, in the sense of Chomsky (2000), is nothing other than the scientific investigation of the world, whatever the aspect involved. Now, scientific investigation of the world, according to Chomsky (1980: 8; 2002), goes along with the

² Cf. Chomsky (2000) and Antony & Hornstein (eds.)(2003) and the interesting reviews of Chomsky (2000) by Stone & Davies (2002), Bilgrami (2002) and Moravcsik (2002).

acceptance of what, using Husserl's expression, the theoretical physicist Steven Weinberg called "the Galilean style" (in physics), i.e. "making abstract mathematical models of the universe to which at least the physicists give a higher degree of reality than they accord the ordinary world of sensation." Whatever aspect of the world is involved, what the scientific (or naturalistic) approach affords is an objective theoretical understanding of the world detached from ordinary human concerns and interests. Because it is based on strong idealizations, this theoretical understanding of the world is bound to be narrow and deep. It is bound to be narrow because the idea of a simultaneous objective theoretical understanding of *all* aspects of the world does not make sense.³ It is bound to be deep because objective theoretical understanding of the world consists in discovering abstract principles which, like fundamental laws of physics, are inaccessible to the resources of human common sense alone, remote from observations and empirical evidence, but from which observations and empirical evidence may be inferred via long chains of explicit reasoning.

Theoretical understanding of the world is not the only kind of understanding accessible to humans. The world also offers us the possibility of artistic (or aesthetic) understanding: "the arts may offer appreciation of the heavens to which astrophysics does not aspire" (Chomsky 2000: 77). But if the goal we are seeking is theoretical understanding of the world, then the idealizations of scientific investigation are not dispensable.

The scientific (or "naturalistic") study of the HLF began in the fifties when Chomsky assigned generative grammar the task of providing an explicit and testable characterization of the computational properties of what is known by any person who is able to produce and understand the sentences of his or her native language.⁴ The task is to describe the recursive procedures which allow the construction of a potentially infinite set of complex linguistic

³ As Chomsky (2000: 69) writes, for example, "the study of communication in the actual world of experience is just the study of the interpreter, but this is not a topic for empirical inquiry [...]: there is no such topic as the study of everything."

⁴ Cf. Chomsky's monumental "Logical Structure of Linguistic Theory" from which Syntactic Structures, published in 1957, was drawn.

expressions out of a finite stock of simple lexical items.⁵ In Chomsky's more recent terminology, mastery of the recursive procedures that allow one to produce and understand a potentially infinite set of sentences from one's native language is an "internal" ("internalized") language or "I-language" and the set of sentences generated by this "I-language" is an "E-language."⁶ An E-language is thus composed of public E-expressions (including sentences) and an I-language is composed of underlying mental I-constructions.

The fundamental task of theoretical linguistics is to understand how an I-language — a stable state of the HLF — allows infinite use of finite lexical resources. Chomsky (1980, 1986) calls this characteristic of the HLF "discrete infinity." In the framework of generative grammar, theoretical understanding resides in computational models of syntactic and semantic processes for constructing complex expressions from elementary constituents.

The goal of generative grammar is to discover the computational properties of the HLF — also called "universal grammar" — based on Chomsky's observation that one's ability to understand and produce sentences from one's native language raises three further questions:⁷

(Q1) What is the system of internal knowledge (I-language) which allows one to understand and produce the sentences of one's E-language?

(Q2) How does this system develop and stabilize in the course of ontogenetic development?

(Q3) How is this system exploited in verbal behavior (both in tasks of production and of understanding)?

Linguistic research shows that the I-language of an adult speaker consists of knowledge (partly explicit, and mainly implicit or tacit) of a vast quantity of syntactic and

⁵ Cf. the contributions by Pollock and by Belletti & Rizzi to the present volume for more detailed discussion.

⁶ "I" stands for "individual", "internal" and "intensional." "E" stands for "external" and "extensional."

⁷ In fact, one could also consider the two further questions: (Q4) How did the human ability to acquire knowledge of some natural language or other arise in the course of phylogenetic evolution? (Q5) How is grammatical knowledge (I-language) of an E-language stored in the human brain?

semantic facts,⁸ including the fact that, in the English sentences (1) and (3), but not (2), the proper name “Mary” can function as antecedent for the pronoun “she” or the possessive adjective “her” of the constituent “her daughter”:

- (1) Mary said that she would come.
- (2) She said that Mary would come.
- (3) Her daughter said that Mary would come.

In answer to (Q1), the task of generative grammar has involved a search for the basic computational principles from which we can infer the fact that in (1) and (3), but not in (2), the noun can function as antecedent for, respectively, the pronoun and the possessive adjective. If an English speaker knows that the anaphor and its antecedent can be bound in (1) and (3), and not in (2), then question (Q2) arises: how does a human child learn this contrast?

The exploration of questions (Q1) and (Q2) was one of the major factors involved in the cognitive revolution that led to the shift from the study of human behavior to the study of the cognitive structures and processes which sometimes result in observable behavior. According to Chomsky, it would be a mistake to think that all interesting questions posed by the use of language can be handled by the scientific (or naturalistic) approach. Chomsky has repeatedly said over the years that the likelihood of reaching theoretical understanding or a scientific explanation of the “creative” aspects of language use is quite low. By contrast, according to him, questions (Q1) and (Q2) are well suited to scientific study. While the investigation of (Q1) is guided by a search for “descriptive adequacy,” the investigation of (Q2) is guided by a search for “explanatory adequacy”⁹: the latter should contribute to the

⁸ Cf. Pollock (this volume), Belletti & Rizzi (this volume) and Chierchia (this volume) for a detailed presentation of these and other relevant syntactic and semantic facts.

⁹ This distinction was developed in Chomsky (1965). Cf. Boeckx & Hornstein (this volume) for clarification of the distinction.

explanation of how the child constructs her I-language (i.e. knowledge of the grammar of her native language) from the primary linguistic data provided by members of her linguistic community. In other words, (Q2) is the question: how do we characterize the initial state of the language faculty by means of which the child converts primary linguistic data into knowledge of the grammar of a particular natural language or E-language?

For forty years, Chomsky has maintained that inspection of primary linguistic data leads to what he calls the “poverty of the stimulus argument,” which in turn is a condition of adequacy on any purported answer to (Q2). This argument involves three complementary premises that could all be summarized by the proposition that the I-language of an adult speaker is vastly underdetermined by the totality of the linguistic data available to a child. Firstly, grammatical knowledge is not the result of explicit learning or teaching: English-speaking parents do not teach their ten-month-old that an English sentence is composed of a noun phrase followed by a verb phrase. Secondly, the utterances that a child encounters are a finite and fragmentary sample of the E-language spoken by members of her community (e.g. English). Thirdly, children acquire knowledge of certain rules for which there are no clues in the set of utterances to which they are exposed: by hypothesis, the corpus of primary linguistic data only contains information of the category “P is a sentence of language L” and no information of the category “P* is *not* a sentence of L.”¹⁰ On the basis of the poverty of the stimulus argument, Chomsky concludes that a child could not acquire knowledge of the grammar of her language unless she was equipped with tacit knowledge of universal grammar, and that this knowledge is a cognitive “module” specialized for the task of language acquisition.¹¹

2. Methodological dualism and common sense concepts

¹⁰ This is the problem of “negative evidence”. Cf. Boeckx & Hornstein (this volume) for detailed discussion.

¹¹ Cf. Goodman (1968), Putnam (1968) and Quine (1969) for objections to Chomsky’s conclusion.

No contemporary philosopher of science would dream of subjecting theories of physics, chemistry, or biology to the authority of a priori conceptual reflection guided by mastery of such ordinary or commonsense concepts as the concepts of *matter, movement, air, fire, vegetable, or life*. Contemporary philosophers of the natural sciences take it for granted that only if theoretical scientific concepts can be freed from the constraints imposed by ordinary commonsense conceptions can scientific theorizing flourish.

Over the years, certain theoretical concepts of generative grammar have been disputed by philosophers of mind and language. In response, Chomsky has made the point that these criticisms tacitly rely on the presupposition that the naturalistic (or scientific) investigation of the HLF can be subjected to some kind of a priori conceptual analysis that accepts the authority of such ordinary commonsense concepts as the concepts of *language, languages, knowledge, mind or mental*. If he is right — as I believe he is —, then he is also right to conclude (Chomsky, 2000: 112) that these philosophical criticisms rest on some kind of intellectual duplicity or double standard: in the natural sciences, the criteria of rationality are based entirely on explanatory success. But the criteria of rationality accepted in the natural sciences are supposed to be simply inapplicable to the study of human cognitive processes, for which the criteria of rationality are supposed to have an entirely independent source. This is the duplicity that Chomsky calls “methodological dualism,” in opposition to methodological naturalism.

The concept of I-language has given rise to two sorts of philosophical perplexity. The first issue is whether an adult speaker of an E-language can truly be said to *know* the grammar of his or her language. A fortiori, can a human baby be said to *know* universal grammar? The second issue is whether the computational explanations of the HLF — or of any other human cognitive capacity — could or should be subjected to conceptual analysis based on the authority of ordinary commonsense conceptions of so-called “mental” phenomena.

When analytical epistemologists wonder whether an adult speaker really knows the grammar of his language, they decompose this question into two sub-questions. First, given a particular E-language (let's say French), does there exist a unique, well-defined set of grammatical rules that generate all and only the sentences of the E-language in question? (This very question involves the contentious presupposition that E-language is given conceptual priority over I-language.) Second, is it appropriate to analyze the cognitive relation between an adult speaker of the E-language and this set of grammatical rules (if it exists) by means of the concept of *knowledge*?

I begin with the first question. In 1963, the philosopher Edmund Gettier published a short article in which he demonstrated that a person S could have a justified true belief in proposition p although S could not be said to “know” that p in the ordinary sense of the word “know.”¹² The majority of analytical epistemologists concluded from this that one ought to give up the traditional idea that having a justified true belief is a *sufficient* condition for knowing a proposition in the ordinary sense. Since 1963, analytical epistemologists have been wondering what other condition should be added in order to turn a true belief that p into genuine knowledge that p. They assume that not unless one truly believes that p can one know that p. They further assume that unless one were introspectively conscious of holding the belief that p — i.e., unless one could express its content verbally by uttering a sentence saying that p and recognize the belief as a belief — one could not believe (truthfully) that p. Since an ordinary speaker is not consciously aware of the grammatical rules of his E-language which he cannot state, the relation between the speaker of an E-language and the grammatical rules cannot be the belief relation, let alone the knowledge relation.

To subject the scientific investigation of the HLF to the constraints of the ordinary concept of *knowledge* is to succumb to methodological dualism. Faced with the theoretical

¹² Cf. Gettier (1963).

successes of molecular biology, philosophers of science would never subject the double helix model of the DNA molecule to the authority of the ordinary concept of *life*. If the theoretical successes of generative grammar do not satisfy the requirements of the ordinary concept of *knowing*, what should we conclude? Only the intellectual duplicity inherent to methodological dualism can block the conclusion that the ordinary concept of *knowing* is inappropriate if we aim to satisfy the demands of the scientific investigation of the HLF.

The first question of analytical epistemologists was whether all and only the sentences of an E-language can rightly be said to be generated by a unique and well-defined set of grammatical rules. To justify a negative answer to this question, Quine (1972) developed an ingenious argument.¹³ Obviously, if it is false that there exists a unique and well-defined set \mathcal{R} of rules which generates all and only the sentences of an E-language, then the question of whether an adult speaker knows \mathcal{R} is irrelevant. To discredit the idea that a speaker has tacit knowledge of a grammatical rule, Quine (1972) offers the distinction between the fact that (verbal) behavior *conforms* to a rule and the fact that it is actually *guided* by a rule. Quine (1972) conceives of the sentences of an E-language on the model of the “well-formed formulas” of the artificial languages of logic. According to him, it cannot be claimed that a single system of rules guides the verbal behavior of a speaker. He relies on two assumptions. First, he supposes that a speaker’s verbal behavior cannot be guided by a rule unless the speaker can formulate and follow the rule consciously. Second, he supposes that, given a set of sentences belonging to an E-language, one can always imagine many rival systems of rules capable of generating the same set of sentences. Quine concludes from this that all a linguist can say that the verbal behavior of a speaker conforms to many extensionally equivalent systems of rules.

¹³ This Quinean argument for the indeterminacy of syntactic hypotheses is independent from Quine’s semantic arguments in favor of his (1960) famous thesis about the indeterminacy of radical translation.

In his answer, Chomsky (2000: 78) observes, on the one hand, that the logical concept of a *well-formed formula* does not apply to natural language sentences. On the other hand, comparison between rival systems of grammatical rules is not limited to extensional equivalence. Since Chomsky (1965), the methodology of generative grammar includes a distinction between weak extensional equivalence and strong intensional equivalence: two rule systems are weakly equivalent if they generate the same set of sentences. Two systems of rules are strongly equivalent if they assign the same structural descriptions to the generated sentences.

In a further move, Quine suggests that the relevant empirical evidence (or observable data) in linguistics is severely limited. He maintains that the evidence in favor of a syntactic or semantic hypothesis about the constituent structure of the sentences of an E-language is strictly limited to the observable verbal behavior of speakers of the E-language. As Quine puts it (1990: 37), “in psychology one may or not be a behaviorist, but in linguistics one has no choice. Each of us learns his language by observing other people’s verbal behavior and having his own faltering verbal behavior observed and reinforced or corrected by others. We depend strictly on overt behavior in observable situations.” Following Quine’s behaviorist assumptions, observing the verbal behavior of monolingual speakers of Japanese surely would not help a French child learn French. Nor could, on Quine’s behaviorist assumptions, monolingual speakers of Japanese enable a French child to learn French by observing, correcting, and reinforcing his verbal behavior.

Clearly, Quine’s argument in favor of the restriction of relevant empirical evidence in linguistics presupposes that a child learning his native language and a linguist are confronted with exactly the same task. But this is a dubious assumption. As Chomsky (2000: 54) emphasizes, acquisition of one’s native language is a largely automatic process in which the child makes no conscious choice. The child applies his initial cognitive capacities (i.e.

universal grammar) to the primary linguistic data made available by members of his community. On the other hand, the linguist makes conscious and laborious use of all possible empirical evidence relevant to discovering the structures, respectively, of the initial innate state of the HLF and the stabilized I-language of an adult speaker.¹⁴ The scientific cues that the linguist can exploit are simply inaccessible to the child.

I will consider just three examples. First, certain discoveries about the neurological structure of the human brain (accessible to the linguist but not the child) are relevant to the comparison among competing linguistic hypotheses. Second, the linguist can systematically compare minimal pairs composed respectively of a sentence of a language and of an ungrammatical sequence composed of the same words. But the ungrammatical sequences constructed by the linguist cannot be part of the primary linguistic data available to the child.¹⁵ Finally, as Chomsky notes (2000: 53-54), examination of what an adult speaker of Japanese knows may indicate that he has tacit knowledge of some abstract syntactic principle P for which no clue exists in the primary linguistic data available to a child learning Japanese. If so, then the generative linguist will have grounds for supposing that this abstract syntactic principle P belongs to universal grammar, the initial state of the HLF. Universal grammar is supposed to be common to children learning Japanese and those learning French. Thus, even if there exist clues for this principle in the primary data accessible to a child learning French, knowledge of this principle by an adult French speaker might derive from universal grammar and not from the fact that he was exposed to utterances in French during childhood. It follows that description of the I-language of an adult speaker of Japanese can be relevant for determining what, in the I-language of a French speaker, is to be attributed to universal grammar and what depends on his personal linguistic experience.

¹⁴ Incidentally, this critique of Quine's analogy between the child's task and the linguist's casts doubt on the version of the "theory-theory" of cognitive development defended by Gopnik (2003). Cf. the response of Chomsky (2003) to Gopnik (2003).

¹⁵ Cf. the argument from the poverty of the stimulus.

Searle (1992) has offered another challenge (distinct from Quine's) to computational explanations of the HLF (or of any other human cognitive capacity). He maintains that any explanation of a genuine mental phenomenon must satisfy the constraint of the so-called "connection" principle according to which a state or process cannot be genuinely mental if its content is not *potentially* accessible to the conscious subjective experience of the human agent to whom it is attributed. The very idea of discrediting computational explanations of human cognitive capacities by glorifying introspection will be judged harshly by those for whom the very task of the cognitive sciences is to unravel the functioning of mental processes whose very existence is inaccessible to mere introspection. Two answers are available in response to Searle's challenge.

On the one hand, the force of the connection principle is weakened by the fact that Searle fails to specify what counts as the *potential* accessibility of some content to the subjective conscious experience of a human agent. What makes the content of some unconscious mental process potentially conscious? Consider a human *blindsight* patient, who has lost the subjective visual experience of the form, contour, size, texture, and color of objects in the part of his visual field affected by a lesion in his primary visual cortex.¹⁶ Should the visual attributes of an object count as *potentially* accessible to the conscious experience of a patient with blindsight on the grounds that they are fully accessible to the consciousness of a healthy subject (without blindsight)? Consider further the phenomenon of so-called "subliminal perception", whereby a healthy human subject is shown a word (from his language) for such a short duration of time that, although he can visually process it, he cannot be aware of it (let alone of having seen it). Nonetheless, it has been shown that, in such conditions, he can extract semantic information carried by the word which, although visually processed unconsciously, facilitates recognition of a second, semantically related, word.¹⁷

¹⁶ Cf. Weiskrantz (1997).

¹⁷ Cf. Marcel (1983).

Should the content of the “subliminal perception” of a stimulus count as potentially conscious on the grounds that, had it been presented slower (for a longer period of time), the subject would have been aware of it? As long as Searle fails to specify what is potentially *inaccessible* to the subjective conscious experience of a human agent, the connection principle runs the risk of being devoid of empirical content, i.e. non-refutable.¹⁸

On the other hand, as Chomsky notes (2000: 75, 106, 134), the connection principle is itself an answer to the question of what the *criterion* (or *mark*) of the *mental* is: on this criterion, accessibility to consciousness is what makes a mental phenomenon genuinely mental. By contrast, no philosopher of the physical sciences believes that he is expected to offer a criterion for what constitutes *mechanical*, *optical*, *electrical* or *chemical* phenomena. But Searle could not apply the connection principle unless he took it for granted that computational theories of the HLF (or of some other human cognitive capacity) can be legitimately subjected to some a priori conceptual reflection based on the authority of the ordinary commonsense concept expressed by the word “mental.” In accordance with methodological naturalism, Chomsky proposes to use the term “mental” on a par with the way physical scientists use such words as “mechanical,” “optical,” “electrical” or “chemical” to refer to different aspects of the world, without presupposing any problematic ontological or metaphysical divisions.

3. Chomsky and metaphysical naturalism

Unlike methodological naturalism, metaphysical naturalism is an *ontological* doctrine. To subscribe to metaphysical naturalism is to subscribe to physicalist monism, which stands in contrast to the ontological dualism between body and mind advocated by Descartes. To subscribe to ontological dualism is to suppose that mental entities are not physical entities,

¹⁸ Cf. Block (1990), Chomsky (2000), Jacob (1995).

because the former are not reducible to the latter. To subscribe to physicalist monism is to assume that all chemical, biological, psychological, linguistic, or cultural processes are physical processes obeying the fundamental laws of physics. For a physicalist, the ontological problem of the relation between body and mind is to find out how to identify the latter with the former via an ontological reduction.

One might have expected Chomsky to appeal to the explanatory success of computational models of the HLF and argue for a physicalist ontology. If mental processes are computational processes, and if computational processes are in turn operations that can be carried out by a machine (built according to the laws of physics), then computational models of a human cognitive capacity show that a machine obeying the laws of physics can carry out operations characteristic of some fundamental human cognitive competence.

This argument has been offered by philosophers who, like Fodor (1975, 1987, 1994), subscribe to metaphysical naturalism. But no trace of it can be found in Chomsky's own work. On the one hand, unlike biological entities (including the HLF), machines, whether abstract or concrete, are artifacts: their functions depend on the intentions of their designers, whose contents in turn raise questions of interpretation that go beyond the limits of scientific investigation. Thus, according to Chomsky (2000: 44-45, 148), understanding the functioning of an artifact cannot contribute to the scientific understanding of any aspect of the natural world — including the HLF, whose structure and function (if any) is independent from the contents of the intentions of any designer.¹⁹ On the other hand, Chomsky could not argue from the explanatory success of computational models of the HLF to physicalism because on his view, for the past three centuries, the ontological controversy between physicalism and substance dualism has been turned into a pseudo-problem between two theses equally devoid of content. This severe diagnosis itself calls for some explanatory comments.

¹⁹ This is an aspect of Chomsky's thesis according to which intentionality-based explanations (of e.g. human action) outrun the limits of naturalistic (or scientific) investigation. I return to this point in section 4.

According to Chomsky (2000: 83-84, 103, 108-109), the problem of the relation between body and mind was a genuine problem in Descartes' time — in the middle of the seventeenth century — when the physical universe was assumed to be governed by the principles of Cartesian mechanics. Since the mind's functioning did not seem to be governed by the laws of Cartesian mechanics, Descartes was rationally led to accept ontological dualism, according to which minds (or mental things) are distinct from bodies (or material things). However, the principles of the mechanical philosophy were swept aside by the explanatory success of Newton's introduction of universal gravitation (a force acting at a distance) as a unifying principle for both celestial and terrestrial mechanics. As Chomsky (2000: 84) puts it, Newtonian mechanics "exorcized the machine," but not the Cartesian conception of the mind.

Arguably, the principles of Cartesian mechanics govern our ordinary or naïve conception of a *physical object*.²⁰ But since it gave up the principles of Cartesian mechanics, theoretical physics no longer has a scientific concept of bodies or physical objects. Chomsky thus throws down a real challenge to physicalists: if the concept expressed by the words "physical object" has no assignable scientific content, then the ontological controversy between physicalism and substance dualism has lost its meaning.

Most contemporary advocates of physicalist monism subscribe to one or another version of the identity thesis between mental and physical entities. Some accept a reductionist version, others a non-reductionist version, of the identity thesis. A few advocate the elimination of mental entities. Because the concept of matter itself has been given up in theoretical physics, Chomsky rejects the very idea (which he takes to be devoid of any scientific content) of any purported physicalist reduction of mental things. Instead, Chomsky advocates the *epistemological* goal of unification (or integration) of scientific theories. But as

²⁰ At least, they seem to govern human infants' "naïve physical" expectations about the behavior of physical objects. Cf. Spelke (1988).

the history of the physical sciences, often referred to by Chomsky (2000, 2002), shows, unification between two scientific theories dealing with entities located at different ontological levels — such as the chemical theory of molecules and the physical theory of atoms — has often required radical modification of the lower-level theory (e.g. the atomic theory) for it to be able to be integrated with a higher-level theory (e.g. the molecular theory).

Chomsky does not exclude the epistemic unification of computational explanations of (higher-level) cognitive competencies with (lower-level) neuroscientific explanations of the contribution of different brain areas to different human cognitive capacities. But the cognitive neurosciences must, he believes, undergo serious revision before their fine-grainedness equals that of computational theories.

Can an advocate of physicalist monism meet Chomsky's challenge by minimizing ontological commitments? For instance, consider the minimal physicalist distinction between two classes of things, both of which are physical, but only one of which turns out to be also mental as a result of its internal structural complexity. Couldn't a physicalist meet Chomsky's challenge by appealing to such a minimal distinction?²¹ The answer is: No, and the reason is that this minimalist version of physicalism still faces Chomsky's (2003: 259) dilemma. Either contemporary physics supplies a complete description of the world or it does not. If it does, then mental entities are simply physical entities (and nothing more). If it does not, then we do not yet know what is (merely) physical.

Chomsky's challenge raises at least three fundamental questions. First, is the epistemological goal of theoretical unification itself really independent from any underlying ontology? Second, what authority should be granted, respectively, to physical theories and to neuroscientific theories in the adjudication of ontological controversies raised by the development of the cognitive sciences? Third, what authority should be granted, respectively,

²¹ Cf. Lycan (2003: 16).

to scientific physical theories and to “naïve” (or common sense) physics in the adjudication of such controversies?

Following Poland (2003), let’s call “methodological physicalism” Chomsky’s view that one should aim for epistemological unification between computational and neuroscientific approaches to the functioning of the human brain. Can methodological physicalism be justified on purely epistemological (or methodological) grounds? If theoretical unity (or simplicity) is a virtue at all, is it a purely epistemic (or even an aesthetic) virtue?

The reason one can (and one ought to) be skeptical is that the epistemic goal that is being sought is unification between theories belonging to different levels, and the very notion of a *level* is an ontological notion. We distinguish the *chemical* level of molecules from the *physical* level of atoms because molecules are things made up of atoms. The neurosciences investigate the structure and functioning of constituents of the human brain. Computational theories investigate so-called “emergent” properties of the brain, such as visual perception or the HLF. Now, Chomsky (2002: 55, 63, 65) treats the view expressed by the cognitive neuroscientist Vernon Mountcastle that “minds, indeed mental things, are emergent properties of brains” as a *truism* (and thus an obvious truth). But it is far from being a truism, since, for example, this is an ontological thesis that serves to justify the choice of methodological physicalism in favor of the epistemological unification between computational and neuroscientific theories of brain processes.²²

Chomsky accepts the ontological assertion that human cognitive capacities are (emergent) properties of the human brain. But he rejects the claim that mental entities are physical entities. An opponent of ontological dualism will conclude that it is an error to subject ontological controversies in the cognitive sciences to the authority of the fundamental concepts of theoretical physics. For an advocate of physicalist monism, the computational

²² An advocate of ontological dualism, e.g., Kripke (1972, 1982), could reject the view that mental properties are emergent properties of the brain.

properties of human cognitive capacities (including the HLF) depend on the neurological structure of the human brain. Even supposing that contemporary basic physical theories of elementary particles are fundamentally incomplete, there is still no reason to believe (as Fodor, 2001 and Lycan, 2003 have noted) that drastic revisions in fundamental theoretical physics would have a major conceptual impact on our scientific understanding of the molecular mechanisms involved in the firing of neurons, the transfer of information between neurons, and the functional roles of distinct brain areas.

Finally, the challenge thrown down by Chomsky to physicalists presupposes that only scientific concepts — not concepts of common sense — can arbitrate respectable ontological controversies.²³ Arguably, what physicalists should do is to reexamine their own conception of the role ascribed to concepts respectively from theoretical physics and from common sense in the adjudication of the ontological controversies underlying the development of the cognitive sciences.

Let us now consider one of the premisses used by Davidson (1970) in support of his own non-reductionist version of physicalist monism, which he calls “Anomal Monism” and according to which any mental event is a physical event, but no psychological concept or predicate is reducible to a physical concept or predicate. According to Davidson, there are, on the one hand, physical laws which subsume the relations between physical events. On the other hand, causal relations may also hold between pairs of mental events and even between pairs of events, one of which is mental and the other is physical. However, according to the premise that Davidson calls “the Anomalism of the Mental,” there are neither genuine psychological laws able to subsume the causal relations between mental events, nor genuine psychophysical laws able to subsume the causal relations between mental events and physical events. On Davidson’s (1970) view, the Anomalism of the Mental (or of naïve psychology) is

²³ According to Chomsky (2000: 139), any scientific investigation, including those that above the level of fundamental physics, creates concepts that have no continuity with the ordinary concepts of common sense.

a reflection of the contrast between what he takes to be the “strict” (i.e. exceptionless) laws of physics and what he takes to be the “truisms” of either psychological or psychophysical correlations.

As Chomsky (2000: 88-89, 138) insightfully observes, the grounds for the anomalism of naïve psychology are also grounds for the anomalism of naïve physics. The gulf between theoretical physics and commonsense or naïve psychology also separates theoretical physics from commonsense or naïve physics. There are no bridge laws linking the concepts of theoretical physics to either those of naïve physics or of naïve psychology. Davidson, who subscribes to a non-reductionist version of physicalism, seems to think that the relation between mental concepts and physical concepts raises a distinctive epistemological problem not raised by the relation between the concepts of naïve physics and the concepts of theoretical physics. But unless Davidson surreptitiously subscribes to some version of methodological dualism, acceptance of the “anomalism of the mental” should force him to accept the anomalism of naïve physics.²⁴

In short, Chomsky seems to be in a position that can be described by means of the following five propositions. First, he emphasizes the fact that there is no room in post-Newtonian theoretical physics for the concept expressed by the term “physical object.” Secondly, on his view, the attempt to unify computational theories of human cognitive capacities with neuroscientific theories of the structure and functioning of the human brain is a legitimate research program. Thirdly, he admits that human cognitive capacities are emergent aspects of the human brain. Fourthly, acceptance of this ontological thesis grounds the search for theoretical unification. Finally, even if there is no room within the scientific

²⁴ On behalf of Davidson, it might be pointed out that the generalizations of naïve psychology raise a problem not raised by the generalizations of naïve physics. Unlike the strict laws of basic theoretical physics, the generalizations of both naïve psychology and naïve physics are *ceteris paribus* and admit of exceptions. But in addition, Davidson (1970) has argued that, unlike the application of concepts of naïve (or commonsense) physics, the application of mental or psychological (either naïve or scientific) concepts to human agents presupposes some normative principle of rationality. Of course, the question is whether Davidson’s arguments in this area presuppose some version of methodological dualism.

theories of fundamental theoretical physics for the concept of *physical object*, still this concept plays a central role in naïve physics, of which the principles of the Cartesian mechanical philosophy are constitutive.²⁵

The concepts of *physical object* and *mental entity* are not part of any scientific theory of the world. The ontological controversy about the nature of the relation between physical objects and mental entities is thus no part of any scientific (or naturalistic) investigation of the world. But the concepts of *physical object* and *mental entity* belong, respectively, to naïve physics and naïve psychology. Chomsky (2000: 90-91) himself calls “ethnoscience” the study of the commonsense conceptual resources by means of which humans, in all cultures, form their stable non-scientific representations of the world. The study of the relations between the concepts of naïve physics and those of naïve psychology thus pertains to naturalistic ethnoscience. If the conceptual representations of the world formed by common sense are features of the human brain, then they are themselves part of the world. If they are part of the world, then ethnoscience is a branch of the (naturalistic) scientific investigation of the world. If so, then the study of the relations between the concepts of *mental entity* and *physical object* is part of the naturalistic investigation of the ordinary conceptual system of representation of the world by human common sense.

4. Chomsky and the naturalization of intentionality

When in the late nineteenth century Brentano (1874) introduced the concept of intentionality into philosophy, he subscribed to the following three theses: first, he took it to be constitutive of intentionality, as it is manifest in such mental acts and states as love, hate, desire, hope, belief, judgment, perception, and many others, that they are directed towards objects that are distinct from themselves. Second, the so-called “objects” towards which the

²⁵ Cf. the work of Spelke (1988) on the cognitive development of naïve physics in human infants. The study of the cognitive capacities of new-born humans reveals the cognitive constraints on the diversity of human culture.

mind is directed by its intentionality have the property that Brentano calls “intentional in-existence.” The notion expressed by Brentano’s word “in-existence” has given rise to much exegetical controversy: did Brentano mean *non*-existence? Did he mean existence *within the mind*? Or both?²⁶ Third, intentionality is the mark of the mental: all and only mental acts and states have intentionality.

Let us assume (in conformity with Brentano’s first thesis) that it is constitutive of intentionality that no one can be said to love, hate, desire, etc. unless there is something that is loved, hated, desired, etc. If this is true, then unless there were *objects* exemplifying the property of *intentional in-existence*, intentionality itself could not be exemplified (second thesis). But love, hatred, admiration, desire and other mental acts are directed not only towards concrete objects but also towards abstract objects (like numbers), mythological constructions (Zeus) or fictional characters (Anna Karenina), which do not exist in space and time. Humans can even entertain thoughts about impossible numerical or geometrical objects such as the greatest prime number or squared circles. Thus, acceptance of Brentano’s first two theses raises a number of fundamental ontological questions in philosophical logic. Are there intentional objects at all? Does recognition of the phenomenon of intentionality force one to postulate the ontological category of intentional objects? Could there be objects that fail to exist? These questions in turn have given rise to a major division within analytic philosophy. The prevailing (or orthodox) response has been a resounding “No.” But an important minority of philosophers, who subscribe to the “intentional objects theory,” have argued for positive responses to these questions. Since intentional objects need not exist, on the intentional objects theory, there are things that do not exist. According to critics of intentional objects theory, there are no such things.²⁷

²⁶ Cf. Jacob (2003).

²⁷ Russell’s theory of descriptions and Quine’s theory of ontological commitment are intended to discredit the theory of intentional objects. But among others, the philosopher Terence Parsons has recently revived the Meinongian theory of non-existent intentional objects. For some detailed discussion, cf. Jacob (2003).

Many contemporary philosophers of mind and language have responded to Brentano's introduction of the concept of intentionality by embracing a so-called *externalist* view of the content of psychological (or mental) states with intentionality. Externalism is the view that intentionality is not an intrinsic property of a cognitive system, but rather a relation between a cognitive system and its environment. In other words, according to externalism, an individual's psychological states (e.g., his or her beliefs, desires or perceptions) derive their contents from the relations holding between the individual (on physicalist assumptions, between the individual's brain) and properties exemplified in his environment. Externalist philosophers of mind and language, therefore, accept the burden of explaining what Soames (1989) (cited by Chomsky, 2000: 132) calls "the fundamental semantic fact of language [...] viz. that it is used to represent the world" — which presupposes that the function of human cognition is to represent the world. To simplify, we can distinguish two versions of externalism: a *normative* and a *descriptive* version: on the former, but not the latter, intentionality is taken to arise from the *norms* obtaining in a linguistic community. Chomsky rejects both versions of externalism.

According to the *normative* version of externalism, the meaning and reference of linguistic expressions of a given E-language are constituted by the linguistic practices accepted by the community of people who speak the E-language in question. On this view, the intrinsic cognitive resources of a speaker are not sufficient to determine the meaning and reference of the words he uses and the beliefs he expresses by uttering them. If some individual did not belong to some linguistic community or another, then his mental states would fail to exhibit intentionality. On the normative version of linguistic externalism, the meaning and reference of the expressions of an external public language (or E-language) are given priority over the meaning and reference of the constructions of an I-language (in

Chomsky's sense). Since they derive from the norms obtaining in a community, meaning and reference are normative properties of public linguistic expressions.

The normative version of externalism is incompatible with Chomsky's methodological naturalism for at least three reasons. First, on Chomsky's view, the expressions of some E-language are the derived products of the underlying psychological constructions that belong to a speaker's internalized I-language. By contrast, the normative version of externalism gives theoretical priority to the notion of a "shared public language" over a speaker's mental representations. Secondly, advocates of normative externalism assume that human verbal communication would simply be impossible unless there existed a single shared external public language, i.e. a set of public linguistic expressions each with a unique public linguistic meaning known to every speaker. Chomsky (2000: 30) rejects the normative externalist picture of verbal communication. He argues instead that verbal communication is a fallible inferential process and that resemblance — not identity — between the external products of different I-languages is sufficient to enable verbal communication.²⁸ Thirdly, Chomsky (2000) constructs numerous examples aimed at discrediting the idea that there exists a fixed well-defined reference relation between the words of a language and non-linguistic entities.²⁹ These examples at least cast doubt on the idea that the reference relation can serve the aims of naturalistic or scientific semantic investigation. I will consider three of these examples:

(4) The bank burned down and then it moved across the street.

(5) The book that he is planning will weigh at least five pounds if he ever writes it.

(6) London is so unhappy, ugly, and polluted that it should be destroyed and rebuilt 100 miles away.

²⁸ This view is also that of Sperber and Wilson (1986).

²⁹ According to Chomsky (2000: 130-31), the Fregean notion of reference (*Bedeutung*) is a technical semantic notion applicable only to the stipulated symbols of Frege's formal language to meet the demands of his logicist project of reducing arithmetic to logic.

Anyone who understands (4) knows both that the implicit subject of the verb “burn” refers to a concrete physical building and that the explicit subject of the verb “move” refers to an abstract institution — which can be physically instantiated as several concrete structures. In this sentence, the phrase “the bank” can, thus, refer at one point to a concrete building, and at another to an abstract institution. In (5), the constituent “the book” is used to refer to an abstract content as grammatical object of the verb “write” and to a physical object as subject of the verb “weigh”. In (6), the anaphoric pronoun “it” and its antecedent “London” are jointly used to refer to agents (its inhabitants), inanimate objects (the buildings), a place and an abstract entity (which can be realized by different concrete entities, of which one can be physically destroyed and the other rebuilt elsewhere).

Chomsky does not use these examples to support the idealist (or non-realist) metaphysical thesis that the world reduces to language and/or mental representations. What these examples are meant to suggest instead is that reference is an *action* carried out by human agents using words which are in themselves devoid of reference. Within one and the same utterance, a pronoun and its antecedent can subtly change reference as a function of the speaker’s *intentions*, without the interlocutor experiencing the slightest difficulty in understanding this change of reference. Given the constant variability of referential intentions, the coordination between speaker and hearer transforms reference into a “mystery.”

Unlike what a human being *knows*, what he *does* is, according to Chomsky, bound to remain a mystery. Generative grammar has shown the way to scientific understanding of an aspect of what a human knows: the human language faculty. But an epistemic divide separates the *problems* encountered in understanding what a human knows and the *mysteries* involved in explaining an intentional action. Although he does not endorse ontological dualism, Chomsky (1980: 79, 1988: 5-6) nonetheless accepts the Cartesian argument for the freedom of the human will, which says that, unlike the behavior of any machine (or mechanical

device), human intentional action is “indeterminate” because a human agent is always free to choose between two distinct courses of action. A human agent can be “incited” to act, but he could always have acted differently. Now, *reference* (i.e. the use of a word to refer to something) is an intentional human action. Thus, because on Chomsky’s view, any act of reference (and what Chomsky also calls the “creative use of language”) involves the freedom of the will, it is presently an epistemic mystery, not a scientific problem.

Most if not all of the philosophers who endorse the *descriptive* (non-normative) version of externalism also subscribe to metaphysical naturalism. Their basic goal is to domesticate intentionality within cognitive scientific psychology, i.e. to show that intentionality can be “naturalized,” or that it obeys the laws of nature. The program of the naturalization of intentionality thus faces two complementary tasks: to show that intentionality has both respectable causes and respectable effects. A philosopher who subscribes to the program of naturalizing intentionality may be tempted to argue, as Fodor (1975, 1987, 1994, 1998) does, from the explanatory successes of computational models of the HLF to the computational-representational theory of the mind (CRTM). But, as I already noted, Chomsky does not. The following four features are distinctive of CRTM.

First, on this conception, all cognitive processes are computational processes. A computational process takes a representation (or symbol) as input and transforms it into a different representation following purely formal rules. CRTM thus presupposes the existence of a “language of thought” (or “mentalese”) composed of symbols that are themselves endowed with syntactic and semantic properties. In virtue of their syntactic properties, the primitive symbols of the language of thought enter into combinations and form complex representations.³⁰ In Fodor’s conception, the primitive symbols (or concepts) of the language of thought are supposed to possess “primitive intentionality,” from which the “derived”

³⁰ A symbol of the language of thought (or concept) is said to be “primitive” if it does not result from the syntactic combination of other symbols (or concepts).

intentionality of all other symbols flows (in particular the derived intentionality of the linguistic symbols of external public languages).³¹

The thesis that the intentionality of symbols in the language of thought has priority over the intentionality of linguistic expressions is of course incompatible with the reversed priority granted to the meaning of public linguistic expressions over the content of mental representations by advocates of the normative version of externalism. Furthermore, the thesis that the intentionality of symbols in the language of thought is prior to that of public linguistic expressions is not immediately subject to Chomsky's objection against the possibility of a scientific understanding of reference. Suppose I hear a dog bark and the processing of this acoustic stimulus causes me to think of a dog. According to CRTM, the perceptual processing of the acoustic stimulus triggers in my language of thought an occurrence of the symbol "Φ" which just is my concept *dog*.

According to CRTM, not all thinking amounts to some intentional (or voluntary) action. For example, the cognitive process whereby my auditory perception of the stimulus is turned into a conceptual representation of a dog is not an intentional (or a voluntary) action. When it results from my perception of a stimulus, my conceptual representation of a dog — i.e. the occurrence of my mental symbol "Φ" — is independent from any intention to refer to a dog.³² If so, then the CRTM approach to the reference of mental symbols is not directly open to the neo-Cartesian argument based on the freedom of the will.

Secondly, according to the version of CRTM defended by Fodor (1994, 1998), the content (or semantic value) of a primitive symbol of the language of thought (i.e. of a primitive concept) results from the nomological correlations between this symbol and exemplifications of a property in the environment.³³ Thus, my primitive concept "Φ" draws its content or semantic value from the fact that it nomically covaries with exemplifications of

³¹ Cf. Jacob (1997).

³² Cf. Jacob (1997).

³³ This follows from the informational semantic view of mental content to which Fodor (1994, 1998) subscribes.

the property of being a dog. In general, the intentionality of a primitive concept derives from psychophysical correlations between a cognitive system and parameters in the environment.³⁴

Thirdly, according to CRTM, psychological explanations are jointly intentional and nomological. The explanation of an action is intentional because what an agent does depends on the content of his intentions, beliefs and desires. It is nomological because the psychological explanation of an action typically consists in subsuming the action under the psychological generalizations that cover the intentions, beliefs and desires of human agents.

Finally, according to CRTM, what makes both psychological explanation *causal* explanation and intentional psychological generalizations *causal* laws is the computational thesis that psychological laws are implemented by underlying computational mechanisms. By virtue of the language of thought hypothesis, the contents of an agent's beliefs and desires reduce to the semantic values of symbols of the language of thought, but the underlying computational mechanisms transform mental symbols solely in virtue of their syntactic properties.

CRTM constitutes the most systematic contemporary effort to assign a role to intentionality in causal psychological explanations and to create a bridge between naïve psychology and computational models in cognitive science. More than anyone else, Chomsky has contributed to promoting computational explanations in cognitive science. But for at least two reasons, he cannot endorse CRTM, which he takes to be a metaphysical project, not a scientific research program. On the one hand, Chomsky accepts the Cartesian argument against the possibility of offering causal explanations of human intentional actions based on the freedom of the human will. On the other hand, on Chomsky's view, purported explanations of human action which rely on the attribution of intentionality to an agent's

³⁴ Fodor (1994, 1998) himself subscribes to an atomistic version of conceptual content according to which the content of every primitive concept is independent from the content of any other primitive concept.

psychological states cannot aspire to the same scientific explanatory value as computational models of human cognitive capacities.

However, even if causal explanations of intentional human actions are at present just proto-scientific theories (not genuine scientific theories), an advocate of CRTM might well object to Chomsky that the Cartesian concept of *freedom* itself belongs to commonsense naïve psychology. So the following question arises for Chomsky: is not the neo-Cartesian argument from freedom of the will an instance of methodological dualism whereby proto-scientific psychological theories of human behavior are subjected to some a priori philosophical reflection guided by the authority of the ordinary commonsense concept of *freedom*?

Because he thinks that the cognitive sciences should emulate the “Galilean style” that has proved so fruitful in the physical sciences, Chomsky takes it that, apart from computational models of a human cognitive capacity (such as the HLF or visual perception), there is presently no serious prospect for a genuine scientific (or naturalistic) investigation of the human mind. As his (2003) replies to Egan (2003) and Rey (2003) clearly reveal, he thinks that CRTM has failed to demonstrate either that intentionality has respectable causes or that it produces respectable effects.

As the history of twentieth century philosophy testifies, the logical and ontological puzzles inherited from Brentano’s intriguing definition of intentionality have given rise to much work in philosophical logic. Furthermore, the evaluation of Brentano’s thesis that intentionality is the mark (or the criterion) of the mental, which has defined much of the landscape of contemporary philosophy of mind, has given rise to many important distinctions, including subtle distinctions between the concepts of *intentionality* and *consciousness*. Finally, some of the gap between the concept of *intentionality* and concepts as widely accepted in the natural sciences as *lawful correlation* and *information* has been filled by the work of philosophers working towards naturalizing intentionality.

It is one thing, however, to make the concept of *intentionality* metaphysically respectable by displaying some of its significant conceptual connections with concepts widely accepted in the natural sciences. (This much contemporary metaphysical naturalism can legitimately claim to have achieved.) It is another thing to show that the concept of intentionality can or should be a constitutive element of computational models of human cognitive capacities that can subsequently give rise to experimental inquiry. Clearly, Chomsky is skeptical about the prospects of a scientific research program into human cognitive capacities that would be based on the concept of intentionality. Furthermore, Chomsky (2003: 274) rejects Fodor's (1987, 1994, 1998) idea that the scientific investigation of human cognitive capacities is continuous with the generalizations of commonsense naïve psychology. The latter, which are intentional, derive from a priori conceptual reflection and are supported by no experimental confirmation. Nor could they be empirically disconfirmed.

In laying out his most recent version of CRTM, Fodor (1998) seems increasingly tempted to endorse a semantic view of computation according to which the preservation of the semantic relations among mental symbols is a constitutive feature of computational processes: "To a first approximation, computations are those causal relations among symbols which reliably respect semantic properties of the relata [...] The essential problem [...] is to explain how thinking manages reliably to preserve truth [...] Turing's account of thought-as-computation showed us how to specify causal relations among mental symbols that are reliably truth-preserving."³⁵ So it may seem as if it is constitutive of a mental representation (or symbol) to which computational processes are applicable that it has content, a semantic value or intentionality. If so, then it might seem as if a computational theory of a cognitive

³⁵ Fodor's (1998) endorsement of a semantic conception of computation is interestingly discussed and criticized by Damian Justo (2007) in his doctoral dissertation, in the context of an evaluation of Fodor's (2001b) arguments against massive modularity. As Justo (ch. 2) rightly points out, it is one thing to recognize that it is a desirable property of a computational system that it preserves the semantic relations among symbols. It is another thing to build this feature into a constitutive feature of computation.

human capacity would be *incomplete* so long as it does not face the question: what do the symbols manipulated by the computational processes represent?³⁶

As Chomsky's (2003) replies to Egan (2003) and to Rey (2003) testify, he would not endorse this semantic conception of computation and would rather keep the definition of a computational process apart from the problems raised by the semantic interpretation of mental symbols. In fact, the doctrine which, in much recent work, Chomsky (2000, 2003) endorses and calls "internalism" seems designed to support just the rejection of such a requirement for determining the content of the mental symbols to which computational processes apply.³⁷ One of the internalist principles he appeals to is a principle of symmetry (or parallelism), which (in accordance with his own minimalist program) governs the computational architecture of the grammars of natural languages. According to this principle of symmetry, the syntax of an I-language generates mental representations on which the rules of phonological and semantic interpretation operate in parallel. The mental representations generated by syntax thus constitute a twofold set of instructions for both the human sensory-motor system (which controls the articulation and perception of the sounds of language) and the human conceptual system (which controls inferences).³⁸

Let M be the mental representation (or I-construction) associated with the proper name "London" by the syntax of an I-language. According to the principle of symmetry, M is a member of two distinct relations involving non-mental entities: on the one hand, M can be thought of as an instruction for the articulatory system enabling the pronunciation of "London." Thus M stands in relation S, the *Sounds* relation, with noises of category N. On the other hand, M is in the purported *Refers* relation R with some presumed non-mental entity E

³⁶ As Peacocke (1994), Egan (2003) and Rey (2003) observe, this question is raised both by computational theories of the HLF and by computational theories of vision.

³⁷ Usberti (2002) develops Chomsky's internalism within the framework of an anti-realist view of meaning partly rooted in Dummett's constructivism and in Meinong's theory of "intentional objects."

³⁸ For a detailed discussion of the theoretical framework of this conception, cf. Belletti & Rizzi (this volume).

(e.g. a city).³⁹ Chomsky (2003: 271) observes that, unless the nature of entities N and E can be well-defined, the relations S and R will remain totally indeterminate. In theoretical linguistics, it is widely taken for granted that a definition of relation S and entities N would be of no scientific interest. By parity, and contrary to most externalist philosophers of language, Chomsky (*ibid.*) concludes that definition of the extrinsic relation R between M and E is scientifically futile.

On the basis of the computational principle of symmetry between the semantic and phonological interpretation of syntactic representations, Chomsky thus distinguishes two notions of representation: a pre-theoretical relational notion and a theoretical non-relational notion. The relational notion is intentionality in Brentano's sense: any representation is a representation *of* something. On Chomsky's view, the pre-theoretical notion can play some auxiliary role in the informal or intuitive presentation of a computational theory, but not within the computational model itself. It is incumbent on the computational theory to explicitly introduce the operational (or formal) notion of representation, which is a purely syntactic notion obeying only the laws of the computational mechanisms. Thus, Chomsky's endorsement of an internalist (or *logical syntactic*) account of computation reveals the extent of the gap that separates him from most philosophers, who, following Brentano's legacy, subscribe to Soames' thesis (1989) that the most basic property of language is that it is being used to represent the world.⁴⁰

References

³⁹ Example (6) was in fact intended to question the existence of the purported relation R.

⁴⁰ Thanks to Richard Carter and Dan Sperber for their comments and also to Richard Carter for helping with the English version.

- Bilgrami, A. (2002) Chomsky and philosophy. *Mind and Language* 17, 3, 290-302.
- Block, N. (1990) Consciousness and accessibility. *Behavioral and Brain Sciences* 13, 4, 596-98.
- Brentano, F. (1874/1911/1973) *Psychology from an Empirical Standpoint*. London: Routledge and Kegan Paul.
- Chomsky, N. (1957) *Syntactic Structures*. The Hague: Mouton.
- Chomsky, N. (1965) *Aspects of the Theory of Syntax*. Cambridge, Mass.: MIT Press.
- Chomsky, N. (1968, 1972) *Language and Mind*. New York: Harcourt Brace Jovanovich.
- Chomsky, N. (1975) *Reflections on Language*. New York: Pantheon.
- Chomsky, N. (1980) *Rules and Representations*. New York: Columbia University Press.
- Chomsky, N. (1982) *The Generative Enterprise*. Dordrecht: Foris.
- Chomsky, N. (1986) *Knowledge of Language, its Nature, Origin and Use*. New York: Praeger
- Chomsky, N. (1988) *Language and Problems of Knowledge*. Cambridge, Mass.: MIT Press.
- Chomsky, N. (1994) Chomsky, Noam. In Guttenplan, S. (ed.) *A Companion to the Philosophy of Mind*. Oxford: Blackwell.
- Chomsky, N. (2000) *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Chomsky, N. (2002) *On Nature and Language*. Cambridge: Cambridge University Press.
- Chomsky (2003) "Replies", in Antony, L.M. & Hornstein, N. (eds.) *Chomsky and his Critics*. Oxford: Blackwell.
- Churchland, P.M. (1993) *A Neurocomputational Perspective*. Cambridge, Mass.: MIT Press.
- Davidson, D. (1970) Mental events. In Davidson, D. (1980) *Essays on Events and Actions*. Oxford: Oxford University Press.
- Davidson, D. (1984) *Essays on Truth and Interpretation*. Oxford: Oxford University Press.
- Egan, F. (2003) Naturalistic inquiry: where does mental representation fit in? In Antony, L.M. & Hornstein, N. (eds.) *Chomsky and his Critics*. Oxford: Blackwell.
- Fodor, J.A. (1975) *The Language of Thought*. New York: Crowell.
- Fodor, J.A. (1987) *Psychosemantics, the Problem of Meaning in the Philosophy of Mind*. Cambridge, Mass: MIT Press.

- Fodor, J.A. (1990) *A Theory of Content and Other Essays*. Cambridge, Mass.: MIT Press.
- Fodor, J.A. (1994) *The Elm and the Expert*. Cambridge, Mass.: MIT Press.
- Fodor, J.A. (1998) *Concepts, Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Fodor, J.A. (2001a) Review of Chomsky's *New Horizons in the Study of Language and Mind*. TLS.
- Fodor, J.A. (2001b) *The Mind Does Not Work That Way*. Cambridge, Mass.: MIT Press.
- Gettier, E. (1963) Is justified true belief knowledge? *Analysis*, 23 121-23.
- Goodman, N. (1968) The epistemological argument. In Searle, J. (ed.) (1971) *The Philosophy of Language*. Oxford: Oxford University Press.
- Gopnik, A. (2003) The theory theory as an alternative to the innateness hypothesis. In Antony, L.M. & Hornstein, N. (eds.) *Chomsky and his Critics*. Oxford: Blackwell.
- Jacob, P. (1995) Consciousness, intentionality and function. What is the right order of explanation? *Philosophy and Phenomenological Research* LV, 1, 195-200.
- Jacob, P. (1997) *What Minds Can Do*. Cambridge: Cambridge University Press.
- Jacob, P. (2003) Intentionality. *Stanford Encyclopaedia of Philosophy*. <http://plato.stanford.edu/>
- Justo, D. (2007) *La modularité de l'esprit revisitée*. Doctoral dissertation, EHESS.
- Kim, J. (1993) *Supervenience and Mind, Selected Philosophical Essays*. Cambridge: Cambridge University Press.
- Kripke, S. (1982) *Naming and Necessity*. Oxford: Blackwell.
- Lycan, (2003) Chomsky on the mind-body problem. In Antony, L.M. & Hornstein, N. (eds.) *Chomsky and his Critics*. Oxford: Blackwell.
- Marcel, A. (1983) Conscious and unconscious perception: experiments on visual masking and word recognition. *Cognitive Psychology* 15, 197-237.
- Miller, G.A. (2003) The cognitive revolution: a historical perspective. *Trends in Cognitive Sciences* 7, 3, 141-44.
- Moravcsik, J.M. (2002) Chomsky's new horizons. *Mind and Language* 17, 3, 303-11.
- Peacocke, C. (1994) Content, computation and externalism. *Mind and Language* 9, 303-35.
- Poland, J. (2003) Chomsky's challenge to physicalism. In Antony, L.M. & Hornstein, N. (eds.) *Chomsky and his Critics*. Oxford: Blackwell.

Putnam, H. (1968) The “innateness hypothesis” and explanatory models in linguistics. In Searle, J. (ed.)(1971) *The Philosophy of Language*. Oxford University Press.

Putnam, H. (1975) The Meaning of “Meaning.” In Putnam, H. (1975) *Philosophical Papers*. Vol. 2, Cambridge, Cambridge University Press.

Quine, W.V.O. (1960) *Word and Object*. Cambridge, Mass.: MIT Press.

Quine, W.V.O. (1969) Linguistics and Philosophy. In Stich, S. (1975) *Innate Ideas*, Berkeley: University of California Press.

Quine, W.V.O. (1972) Methodological reflections on linguistic theory. In Davidson, D. & Harman, G. (eds.)(1972) *Semantics of Natural Language*. Dordrecht: Reidel.

Quine, W.V.O. (1990) *Pursuit of Truth*. Cambridge, Mass.: Harvard University Press.

Rey, G. (2002) Chomsky, intentionality and a CRTT. In Antony, L.M. & Hornstein, N. (eds.) *Chomsky and his Critics*. Oxford: Blackwell.

Searle, J. (1992) *The Rediscovery of the Mind*. Cambridge, Mass.: MIT Press.

Soames, S. (1989) Semantics and semantic competence. *Philosophical Perspectives* 3, 575-96.

Spelke, E. S. (1988) The origins of physical knowledge. In Weiskrantz, L. (ed.)(1988) *Thought Without Language*. Oxford: Oxford University Press.

Sperber, D. & Wilson, D. (1986) *Relevance, Communication and Cognition*. Cambridge: Mass. Harvard University Press.

Stone, T. & Davies, M. (2002) Chomsky amongst the philosophers. *Mind and Language* 17, 3, 276-289

Usberti, G. (2002) Internalism and anti-realism: a proposal. Mimeo.

Weiskrantz, L. (1997) *Consciousness Lost and Found*. Oxford: Oxford University Press.