

Noam Chomsky, *What Kind of Creatures Are We?* Columbia University Press, 2016, 200pp, \$19.95 (hardcover), ISBN 978-0-231-17596-8

Reviewed by Pierre Jacob, CNRS, Institut Jean Nicod, Ecole Normale Supérieure, Paris, France.

Noam Chomsky's slim volume is comprised of four chapters (the last of which was published earlier by the *Journal of Philosophy*¹), preceded by an excellent Introduction by the philosopher Akeel Bilgrami. The broad question considered by Chomsky gives the book its title: *What Kinds of Creatures Are We* (henceforth abbreviated as WKC)? Chomsky breaks it down into three more specific questions: What is human language? What are the limits (if any) of human understanding? What is the common good to which we should strive?

Chomsky's treatment of the first two questions combines tools from generative linguistics, the philosophy of the cognitive sciences, the history of philosophy and the history of science. For lack of space, I shall restrict myself to the first two questions, focusing on three main topics: Chomsky's approach to the evolutionary origins of the human language faculty, his distinction between problems and mysteries, and his approach to the relations between mind and brain.

In the 1950's Chomsky laid the foundations for a new scientific approach to the human language faculty based on the concept of *generative grammar*. A generative grammar is an explicit characterization of the tacit knowledge that enables a speaker to produce and understand any sentence of her native language. A sentence is an ordered sequence of words

¹ Cf. Chomsky, N. (2009) The mysteries of nature: how deeply hidden? *The Journal of Philosophy*, cvi, 4, 167-200.

and morphemes that belong to the lexicon of the language. As Chomsky has argued, there is no upper grammatical limit on the length of a sentence of a natural language. If so, then while the lexicon is a finite list, the set of sentences that comprise a natural language is not finite. Thus, a generative grammar must have the *recursive* resources necessary to generate a non-finite set of sentences out of a finite set of lexical items.

Moreover, Chomsky drew a methodological distinction between *weak* and *strong* generative capacity: a grammar weakly generates an infinite set of linearly ordered strings and strongly generates an infinite set of hierarchically structured constituents, each of which can be mapped onto a linear string. What is primary is *strong*, not weak, generative capacity: the proper object of syntactic inquiry is not merely to offer one of many possible *extensionally* correct derivations of an infinite set of strings (an E-language), but instead an I-language, i.e. the correct *intensional* characterization of the internal recursive procedure that is mentally represented in a speaker's mind.

As Chomsky has pointed out, it is one thing to offer an intensional account of what a mature speaker of a natural language tacitly knows. It is another thing to address the *ontogenetic* problem of language acquisition in light of what Chomsky called "the poverty of the stimulus." How does a healthy human infant come to acquire the tacit knowledge of the grammar of her native tongue, on the basis of her access to primary linguistic data (i.e. the linguistic evidence made available to her by members of her linguistic community)? Chomsky argued that, given the *poverty* of the stimulus, normal language acquisition would be a miracle if a healthy human child did not come genetically equipped to the task with tacit knowledge of a set of principles, which he called "universal grammar" (UG) and which could impose narrow constraints on the grammatical form of possible human languages. To the extent that language acquisition is uniquely human and not based on explicit instruction, UG can be hypothesized to be species specific and part of human biological endowment.

Much linguistic research in the framework of generative grammar has been driven by the goal of resolving the tension between two desiderata, which Chomsky (1965) called respectively *descriptive adequacy* and *explanatory adequacy*.² On the one hand, particular grammatical rule systems must reflect the known diversity of human languages and be descriptively adequate. On the other hand, the diversity of rule systems must be limited so as to be consistent with (or derivable from) the narrow constraints imposed by UG and thereby meet the challenge of explanatory adequacy, i.e. contribute to explaining language acquisition. In the early 1980's and subsequent years, for the purpose of mitigating this tension, Chomsky adopted the so-called "Principles and Parameters" (P&P) approach, which assumes that UG is linked to a "switch-box" of (a finite list of) parameters, each of which can be set to one of a pair of values. The child's task is to discover which particular setting of the parameters determines her native tongue, on the basis of her ontogenetic experience.

In recent years, pushing further the quest towards radical simplification of the computational properties of UG, Chomsky has endorsed "the strong minimalist thesis" (SMT). The driving force behind P&P was to shed light on the *ontogenetic* problem of language acquisition. One of the further goals of SMT is to also address the *phylogenetic* problem of the evolutionary origins of UG's *Basic Property*, namely to "provide an unbounded array of hierarchically structured expressions" (WKC, p. 4).

In WKC (pp. 41, 125), Chomsky suggests that this phylogenetic question can be further divided into two separable questions: (1) what are the evolutionary origins of the core computational procedures of UG? (2) What are the evolutionary origins of "the atoms of computations"? While SMT offers an intricate set of scientific grounds for addressing (1), much, if not most, of Chomsky's examination of (2) in WKC rests on interesting criticisms of referentialist (or externalist) approaches to the semantics of atoms of computations.

² Cf. Chomsky, N. (1965) *Aspects of the Theory of Syntax*, Cambridge, Mass., MIT Press.

It is part of SMT that the output of the core syntactic operations of UG is made available to (or interpretable by) two distinct interfaces, the first of which is the *sensorimotor* interface in charge of the process of *externalization*, i.e. communication. The second *semantic-pragmatic* interface is involved in (mostly internal) *thought* processes. A further condition of adequacy on any purported answer to (1) is that it should reduce the complexities of the syntactic computational procedure to the simplest possible recursive operation(s). What makes Chomsky's SMT-based answer to (1) tantalizing is that, if correct in broad outline, it vindicates a deep asymmetry between the ways the core syntactic component of UG relates to the sensorimotor and to the thought interface. In a nutshell, it resurrects an intriguing version of the language of thought hypothesis.

SMT involves the strong hypothesis that all syntactic complexities reduce to an operation called Merge, which takes two objects X and Y and constructs the new object Z, $\text{Merge}(X,Y) = \{X,Y\}$, in accordance with the principle of *Minimal Computation*, which stipulates that “neither X nor Y is modified by Merge, and that they appear in Z unordered” (WKC, p. 16). If neither X nor Y is part of the other, then *External Merge* applies: for example, the combination of $X = \textit{read}$ and $Y = \textit{that book}$ generates *read that book*. If one is part of the other, then *Internal Merge* applies, as illustrated by the combination of $X = \textit{John read which book}$ with $Y = \textit{which book}$, which generates *which book John read which book*, which Chomsky says “surfaces as ‘which book did John read’ by further operations” (WKC, p. 17).

The sensorimotor interface involves the collaboration between a *motor* and a *perceptual* component. The motor component is subject to a Principle of Minimal Computation: “compute and articulate as little as possible” (WKC, p. 19). As a result, it articulates *only* one of the two copies of the constituent *which book* (both of which are present in the output of Internal Merge), which in turn makes the task of the perceptual component of

externalization (and communication) correspondingly costly. There is one pending issue that, so far as I can see, Chomsky does not explicitly address in WKC. On the one hand, it follows from the fact that Merge is subject to the Principle of Minimal Computation that X and Y are unordered in the set $\{X,Y\}$. On the other hand, it seems as if $Y = \textit{which book}$ directly lands to the left of $X = \textit{John read which book}$ in the structure yielded by Internal Merge: *which book John read which book*. So the question is: should the direct output of Merge, which is made available to the semantic-pragmatic interface alone, be the unordered set $\{X,Y\}$? Or should it instead be the ordered pair $\langle X,Y \rangle$, in accordance with further constraints of linear order imposed by the sensorimotor interface?

Bracketing momentarily this intriguing issue, the purported reduction of the syntactic complexities to Merge has several potential implications, one of which is that the phenomenon of *displacement* (ubiquitous in natural languages, but not found in formal artificial languages) turns out to be a direct consequence of Merge. In the sentence “Which book did John read?” the phrase “which book” can be heard at one place and interpreted somewhere else (WKC, p. 17). This in turn highlights one of the many conflicts between communicative efficiency and the principle of Minimal Computation. Chomsky argues that whenever there is a conflict, the latter trumps the former.

Chomsky further stresses two basic differences between the semantic-pragmatic and the sensorimotor interfaces. First, while core syntax is linked to a single semantic-pragmatic interface, it cannot be linked to a single sensorimotor interface. Spoken language recruits the motoric resources of the vocal system and the perceptual resources of the auditory system. But sign language recruits the motoric resources of hand movements and the perceptual resources of the visual system. Secondly, while the sensorimotor system tracks linear order, the semantic-pragmatic system alone is sensitive to the long distance hierarchical dependencies created by syntax.

The picture of the human language faculty that arises from SMT can be summarized by the joint conclusions that “externalization of language... is an ancillary process” (WKC, p. 125) and “language is an instrument of thought” (p. 14). At one point, Chomsky even writes that “externalization is rarely used... most use of language by far is never externalized... it is a kind internal dialogue” (p. 14). This raises the mind-boggling possibility that human thought processes might be so disembodied as to enable one to directly represent the semantic content of some complex array of hierarchically structured expressions without subjecting it to the constraints of linear order. But if “internal dialogue” means silent speech, then thought may require linear order via motor imagery, i.e. the off-line use of the sensorimotor interface.

Chomsky’s purported answer to (1) is that some time around 100.000 years ago or less, “some slight rewiring of the brain” of one individual, caused by a single genetic mutation, “yielded Merge... providing the basis for unbounded and creative thought” (p. 40). This mutation must have afforded the individual entirely new powers for *thought*, not for communication, provided it occurred in the genotype of a single individual who could transmit it to his or her progeny, not use it right away for communication with conspecifics.

Turning to phylogenetic question (2), we shift to the evolutionary origins of the *atoms* of computation, whose fundamental semantic nature is puzzling for at least two reasons. First, *atomic concepts* are “word-like objects, but not words”: unlike lexical items, which are involved in sensorimotor computations, the atoms of syntactic computations lack phonological properties and reach the conceptual-intentional interface only (p. 41). Secondly, Chomsky strongly objects to *referentialist* (or externalist) approaches to the contents of atomic concepts. Along with Peter Strawson, he assumes that speakers, not words, can perform referential actions with their use of words: unlike symbols from formal artificial languages and signs from animal communication systems, words of natural languages by themselves do not refer to extra mental entities. For example, ordinary words, e.g. ‘house’

(just like names of cities or persons), do not straightforwardly refer to non-mental entities: they can be used to refer to concrete objects with mind-dependent properties that can be physically destroyed and rebuilt elsewhere. Furthermore, whether a set of objects (e.g. tree branches) is a *thing* may depend on the presence of human (artistic) purposes and intentions (WKC, pp. 42-46, 126).

There is a subtle link between Chomsky's rejection of referentialism and his endorsement of *mysterianism*. If all signs from non-human animal communication systems, but not all human atomic concepts, can be directly mapped onto mind-independent events and properties, then there is no homologous counterpart to human atomic concepts in "the animal world." If so, then human scientific inquiry may be deprived of access to empirical evidence relevant to answering phylogenetic question (2) about the evolutionary origins of the atoms of computation (WKC, pp. 41-42, 48). Clearly, Chomsky's claim about the potential limits on human access to empirical evidence should not be construed as the claim that the conceptual issues raised by the evolutionary origins of the atoms of computation fall beyond human cognitive capacities. It nonetheless paves the way for Chomsky's renewed defense of mysterianism, which rests on his earlier distinction between "*problems*, which fall within our cognitive capacities, and *mysteries*, which do not" (WKC, p. 27).

Chomsky's mysterianist distinction between problems and mysteries attractively reflects his deep commitment to methodological naturalism and his correlative rejection of methodological dualism. Just like the physical and cognitive capacities of members of other biological species, the physical and cognitive capacities of humans have both scope and limits. Furthermore, they would not have scope unless they had limits. Chomsky's biolinguistic approach to the human language faculty is a perfect illustration: UG is part of human biological endowment. It specifies the class of languages accessible to humans only. It underlies a speaker's tacit knowledge of a finite system of grammatical rules, which in turn

enables her to produce and understand an infinite set of sentences. “But infinite is... not the same as limitless. English is infinite, but doesn’t include Greek. The integers are an infinite set but do not include the reals” (WKC, p. 55).

However, it is one thing to posit the *existence* of some mysteries-for-humans. It is another thing to specify the contents of any particular mystery-for-humans, a fortiori to offer a principled characterization of the full class of mysteries-for-humans. To further spell out the contents of specific mysteries-for-humans, Chomsky’s strategy involves four ingredients: he turns to the history of science and philosophy, in particular to the historical investigation of the seventeenth century scientific revolution. He posits a distinction between *common sense* understanding of the world and *scientific* investigation. He posits a distinction between common sense understanding of the *physical* and the *mental* aspects of the world. Finally, he posits the existence of what he calls the human “science-forming faculty.”

On Chomsky’s picture of the scientific revolution of the seventeenth century, Newton’s introduction of action at a distance destroyed the mechanical philosophy to which Newton and his contemporaries remained deeply committed. What this shows is that while the principles of the mechanical philosophy faithfully reflect *common sense* understanding of physical aspects of the world, *scientific* inquiry can and must depart from common sense understanding. While *action at a distance* is a mystery for human common sense understanding of physical aspects of the world, following Hume’s precept, scientific inquiry should adopt “mitigated skepticism,” i.e. give up the utopian epistemic ideal that “the world will be intelligible to us” (WKC, pp. 89-90, 105).

However, the quest for scientific explanation of *mental* aspects of the world may be unable to overcome and disregard some of the mysteries of common sense understanding if not for ever, at least for a long time to come. For years, following Descartes, Chomsky has urged that the creative use of language, i.e. the ability to use language in innovative ways that

are appropriate to, but not caused by, circumstances, is a mystery-for-humans. It is in fact a particular instance of the mystery of free will, whereby humans feel free to act and never compelled to do so. Of course, humans might spontaneously form the intuitive belief that they are free to act while they may not be free to so believe. But Chomsky himself seems reluctant to explore this possibility: he presumably thinks that so far neuroscientific research into human decision-making lacks the experimental and the theoretical resources to overcome and disregard one particular mystery of common sense understanding of the *mental* aspects of the world, namely the human feeling of free will (WKC, pp. 94-96).

One last intriguing feature of Chomsky's mysterianism lies in his commitment to a hypothetical "science-forming faculty" (SFF), which "provides [the human mind] with a limited array of 'admissible hypotheses' that are the foundations of human scientific inquiry" (WKC, p. 28). Clearly, Chomsky assumes that like other biological systems (e.g. rats' ability to run mazes), SFF has its potential scope and limits. In other words, SFF stands to the class of scientific hypotheses accessible to humans in the same relation as UG stands to the class of languages accessible to humans and for the same reasons (WKC, pp. 104-105): the human ability for scientific investigation would not have scope unless it had limits. If so, then Chomsky's distinction between problems and mysteries should apply to SFF.

Chomsky's commitment to SFF raises at least two puzzles for his own framework. First of all, the analogy between SFF and UG is open to doubt. As Chomsky has pointed out, language-acquisition cannot rest on explicit teaching. But learning the content of a scientific theory does rest on explicit teaching. Moreover, unlike an individual's internal thought processes, scientific inquiry involves social cooperation and public discussion and disagreement among various individuals. Secondly, there is a potential tension between Chomsky's view that the distinction between problems and mysteries applies to SFF and his recommendation that scientific investigation is (and should be) guided by "mitigated

skepticism.” To adopt mitigated skepticism is to *dismiss* the utopian idea that “the world will be intelligible to us” (WKC, p. 90). To the extent that scientific inquiry, not commonsense, is guided by mitigated skepticism, unlike commonsense, scientific inquiry should disregard, not “stare in puzzlement” at, the mysteries of commonsense.³

On Chomsky’s view, the *ontological* version of the mind-body problem makes good sense *only* against the background of the principles of the mechanical philosophy, which lie behind human common sense understanding of the basic properties of *physical* aspects of the world. Descartes’ unrestricted acceptance of the mechanical philosophy made ontological dualism inevitable. Newton’s introduction of action at a distance destroyed the principles of mechanical philosophy and turned the very notion of a material object (or body) into a mystery for common sense understanding. It thereby rendered the ontological version of the mind-body problem obsolete (if not meaningless). But as Chomsky’s intricate historical discussion in WKC further highlights, Locke, Hume and especially Priestley soon interpreted the scientific rejection of the principles of the mechanical philosophy (by Newton’s introduction of action at a distance) as grounds for dismissing Cartesian ontological dualism and for embracing the loose physicalist notion of “thinking matter.” They were thereby advocating early versions of the prevalent contemporary view that minds are “emergent properties of brains.”

Unlike many contemporary philosophers of mind, Chomsky thinks that the traditional ontological mind-body problem has evaporated. He further thinks that we presently face interesting epistemological issues of *unification* among theories of various aspects of the world. In particular, we face the issue of the unification between computational approaches to human cognition (including the human language faculty) and experimental neuroscientific investigations of the constituents, the mechanisms and the organization of the human brain

³ Cf. Chomsky, N. (2000) *New Horizons in the Study of Language and Mind*, Cambridge, Cambridge University Press, p. 145.

(and other brains). Unlike many neuroscientists, Chomsky does not expect *reduction* of computational models of human cognition to current neuroscientific theories of the brain. Clearly, chemistry deals with larger entities than fundamental physics: in some obvious sense, the properties of molecules depend on the properties of their atomic and sub-atomic constituents. Nonetheless, as Chomsky provocatively points out, the unification between chemistry and physics had to wait until physics, not chemistry, underwent radical changes in the twentieth century.

Even though “contemporary neuroscience is hardly as well established as physics was a century ago” (WKC, p. 36), there are parallels to be drawn with respect to the purported unification between neuroscientific knowledge of the brain and the cognitive scientific approach to the human language faculty. Clearly, neuroscience, not theoretical linguistics, can provide experimental knowledge about the firing of neurons (i.e. elementary constituents of the brain) based on e.g. single cell recording. However, along with vision scientist David Marr and cognitive neuroscientist Randy Gallistel, Chomsky assumes that what primarily matters for the purpose of unification is that the neuroscientific investigation of brain mechanisms should attend to the tentative answers to two fundamental preliminary questions: what is the computational task performed by the relevant cognitive system (whether UG or vision)? What algorithms carry out the computations? While linguistics might never reduce to neuroscience, unification may be delayed until neuroscientific models of the brain, not computational models of the human language faculty, undergo radical changes.