Two notions of modularity

Jerry Fodor opens his deservedly influential monograph on modularity by recalling that the butterflies were fluttering in a joint seminar that we taught in 1980. There were actually two highlights of the seminar: Jerry's ideas about modularity, which grew into the caterpillars presented in the monograph, and his early thoughts about the intriguing conceptual atomism that he was developing at the same time. I benefitted from lively discussions with him about both – with a mixture of puzzlement, accord, disagreement. These remain, but can be sharpened in the light of progress since. I will keep here to the issue of modularity.

The monograph was inspired by an observation about perception of speech by Merrill Garrett: that parsing is "basically…a reflex." Correspondingly, the focus is on input systems and fixation of belief in central systems. The major example is parsing in language, with some observations about vision and other input systems.

Fodor briefly alludes to a different process that is "basically…a reflex" (p. 100): acquisition of language: "the neural mechanisms subserving input analysis develop according to specific, endogenously determined patterns under the impact of environmental releasers." With one important qualification, that has been a guiding thesis of the study of generative grammar within what has come to be called "the biolinguistic framework" and "the generative enterprise" since its origins in the 1950s.

The qualification is that these neural mechanisms are not limited to "subserving input analysis" (parsing). They also subserve production: the normal use of language to express thoughts either externally or in "internal dialogue."[1] In this crucial respect, language is quite different from vision and the other input systems that Fodor discusses. Correspondingly, it is not clear that language falls within Fodor's framework, let alone that it can serve as the major illustration.

It is perhaps worth noting that it is the production of expressions, not parsing, that has been the central concern of scientific-philosophical inquiry into language since the early days of the modern scientific revolution. Galileo and the Port Royal logicians and grammarians were awed by the "marvelous invention" of a means to construct "from 25 or 30 sounds that infinity of expressions, which bear no resemblance to what takes place in our minds, yet enable us to reveal [to others] everything that we think, and all the various movements of our soul." Galileo regarded that as an achievement "surpassing all stupendous inventions," even those of "a Michelangelo, a Raphael, or a Titian." For Descartes, this was a primary difference between humans and any beast-machine and provided a basis for his mind-body dualism. Wilhelm von Humboldt conceived language to be "a productive activity" that makes "articulated sound capable of expressing thought" – "audible signs for thought," in the words of William Dwight Whitney. For the last great representative of this tradition, Otto Jespersen, the central question of the study of language is how its structures "come into existence in the mind of a speaker" on the basis of finite experience, yielding a "notion of structure" that is "definite enough to guide him in framing sentences of his own," crucially "free expressions" that are typically new to speaker and hearer.

[1] Statistically, by far the major use of language. There is reason to suspect that most of it is inaccessible to consciousness. For some comments, see Chomsky (2013b,c).

EFTA01203404

In contrast, the input aspect of language use does not seem to have been a major concern.

Evidently, the input and output (production) systems are linked. No one understands only Japanese and speaks only Swahili. The natural assumption -- which, to be clear, I've always assumed to be correct -- is that language is a module of a "central system," which is accessed in the many kinds of use of language, including input analysis and externalization in production.

Like Fodor, I will keep here to speech. Externalization, when it takes place, along with parsing, appears to be independent of sensory modality, a matter of some importance that I will put aside, though it bears directly on a question addressed below: How significant are input modules for inquiry into the nature of human language?

The tradition does not distinguish production from generation, the latter akin to the process whereby an axiom system generates proofs and theorems, more generally the way a finite program can determine an infinite array of symbolic objects. That distinction, which is crucial, had become quite clear by mid-twentieth century, thanks to the work of Gödel, Turing, Church and others, making it possible to capture more clearly and to pursue intensively some of the ideas that animated the tradition. The postulated acquisition-based module of the central system is a generative process accessed for production and input analysis. It is what has come to be called an I-language (internal, individual, intensional), in earlier years called "a grammar" in one of the uses of this systematically ambiguous term.[2]

It is also worth emphasis that in much of the tradition, the aspect of production that was salient was its free creative character, the constantly innovative use of "free expression" in ways that are appropriate to circumstances but apparently not caused by them, and that elicits thoughts in the hearers that they can them formulate themselves. That aspect of human action, for Descartes most vividly revealed in language use, remains at the border of scientific inquiry, or beyond, a fact recognized in the most sophisticated studies of voluntary action. As the point is put ("fancifully") by Emilio Bizzi and Robert Ajemian, "we have some idea as to the intricate design of the puppet and the puppet strings, but we lack insight into the mind of the puppeteer."[3] Similarly in the case of language, in terms of Humboldt's now often-quoted aphorism that language involves "infinite use of finite means," we now have some idea as to the nature of the means that are used, but the mind of the user remains a mystery. The problem reaches beyond Fodor's "First Law of the Non-existence of Cognitive Science" (107).

The two types of modules suggest two different ways of approaching the nature of language: as a parsing system (with production and the nature of the linkage left to the side) or as an internal cognitive system accessed by various uses of language. The former approach is sometimes held to be the only one possible, with human language taken to be "by definition an experience-dependent mapping of auditory and visual stimuli onto meaningful objects, actions, and concepts."[4] The traditional perspective, and the development of certain aspects of it in the

[2] The term "I-language" was introduced in Chomsky (1986), after Fodor's book appeared. The purpose of the terminological change was to overcome the ambiguity, which had often been misleading, and to clarify what was meant by "grammar" in the relevant sense.
[3] Bizzi and Ajemian (2015).
[4] Albright (2015).

EFTA01203405

biolinguistic-generative framework, is, however, at least an alternative, and one that I think is more revealing of the nature of language, for reasons to which I will return.

Central-system modularity may appear to be inconsistent with Fodor's rejection of modularity for central systems, but it actually is not. Let us put that problem aside for the moment, just assuming consistency.

One major contribution to the biolinguistic program was Eric Lenneberg's fundamental work (1967), which founded modern biology of language and along with much else, formulated the basic issues of evolution of language with clarity and insight that remain unsurpassed, and also provided important evidence about dissociation of language from other cognitive faculties. The latter topic has been extensively pursued since, yielding the conclusion that the language faculty is "a distinct module of the mind/brain, based on domain-specific organizing principles,"[5] and accordingly lending support to the thesis that acquisition-based modularity of central systems is a real phenomenon. It is this concept of central-system modularity that is developed, for example, in Chomsky (1975, chap. 1) and many other publications before and since. This "modular view of learning" is "the norm in neuroscience" today, Randy Gallistel observes, referring in particular to the "module for learning language." It is appropriately called an "organ," he continues, because "the specialization of structure and function that we see in organs is an appropriate way to think about the specialization of structure and function we see in the various learning organs." In general, learning is based on specialized mechanisms, "instincts to learn" in specific ways, yielding modules within the brain that perform specific kinds of computation such as in the remarkable navigational feats and communication capacities of insects. Apart from "extremely hostile environments," these modules develop and change states under the triggering and shaping effect of external factors, more or less reflexively, and in accordance with internal design. That is the "process of learning,"[6] though "growth" might be an appropriate term.

In this respect, language acquisition falls together with vision and other input systems of the mind-brain – though language crucially differs from them in that it provides not only an "input system" but also an output system and, I presume, a central generative system that both access.

Uncontroversially, the systems involved in navigation, vision, and other subcomponents of the organism ("modules") are in substantial part genetically determined. For language, the theory of the genetic component has been called "universal grammar (UG)" in contemporary work, adapting a traditional term to a new framework.

Curiously, though adopted without serious question elsewhere, the assumption for language is considered highly contentious, if not refuted by "field linguists."[7] In one formulation, "Universal Grammar is dead."[8] The only coherent interpretation of this thesis is that language is acquired by other cognitive capacities that are somehow unique to humans. The suggestion faces two problems: one is the failure to deal with even the simplest cases, such as those discussed below.

---

[5] See Curtiss (2013) for review of a wide variety of evidence.

[6] Gallistel (1999a,b).

[7] For example Churchland (2013). In fact, "field linguists" – that is, linguists who work with the wide variety of languages that have come under investigation since the early days of generative grammar – have repeatedly demonstrated the opposite: that languages that appear to vary widely on the surface are in fact cast to much the same mold when investigated in depth.

[8] Tomasello (2009). His comments suggest that he may be misinterpreting UG in the manner discussed in note 9.

EFTA01203406

The other is the radical (double) dissociations that have been found since Lenneberg's pioneering work in the 1950s. I will put these beliefs aside here.[9]

The central systems incorporate principles, which enter into behavior. For bee communication, for example, internal principles enable calculation of the changing angle between the sun's azimuth and the direction of the food. We assume these to be neurally coded in some manner, though how is apparently not well understood. For vision and language we find such principles as (1), (2), respectively:

(1) The Rigidity Rule
(2) The Rule of Structure-dependence[10]

In these and other cases, the crucial question is why the principles hold.

The Rigidity Rule, as defined by Donald Hoffman (1998) in his study of visual intelligence, holds that when other rules permit, image projections are interpreted "as projections of rigid motions in three dimensions," even with highly impoverished stimuli. That seems initially problematic. The environment throughout the evolution of the visual system rarely contained rigid objects, and the experimental work on the principle shows that presentations are perceived falsely. Questions thus arise about the internal nature of the visual system, the factors in its development in the individual, and its evolution. Related questions were raised by Descartes in his work on the visual system, for example, when he speculates (plausibly) that presented with the drawing of a triangle, a child will not take it to be the "composite figure of the triangle drawn on paper…but rather the true triangle," because "the idea of the true triangle was already in us," as an innate concept. In Ralph Cudworth's formulation, the intelligible idea of an object is not "stamped or impressed upon the soul from without, but upon occasion of the sensible idea excited and exerted from the inward active and comprehensive power of the intellect itself," based on its innate structure, a version of the idea that experience conforms to the modes of cognition.

One of many illustrations of case (2), the principle of structure-dependence, is given by (3)-(6):

(3) Birds that fly instinctively swim
(4) The desire to fly instinctively appeals to children
(5) Instinctively, birds that fly swim

---

[9] Sometimes these beliefs are based on confusion between UG and "language universals," that is, properties found quite generally in language, like Greenberg's famous universals. Such generalizations are, of course, expected to have exceptions, which, like the generalizations themselves, are a valuable stimulus to research. Another common claim is that UG consists of only "tendencies," which, if there were any reason to believe it, would leave us in an even more difficult position: what is the genetic basis for the "tendencies"? Fodor uses the term "linguistic universals" in the sense of UG, but he was writing before the current confusions infected the fields.

[10] This is the one case that has been subjected to extensive efforts to account for the facts by general learning mechanisms. All efforts are irremediable failures (Berwick et al., 2011), though the failure is in fact much deeper than discussed there: the wrong question is being addressed. The right question is *why* the principle holds for all constructions in all languages. The methods proposed would work just as well for a linguistic system in which the simpler linear computation held. The studies keep to the case of auxiliary inversion, a limitation that suggests (erroneously) that adequate data might be available to the child. The illusion is quickly dispelled by construal examples such as (3)-(6) below. One common fallacy is that the results follow from the fact that hierarchy is available – as is linear order, in fact far more saliently in presented data.

EFTA01203407

(6) Instinctively, the desire to fly appeals to children.

The structures of (5) and (6) are, roughly, as indicated by bracketing in (5') and (6') respectively:

(5') Instinctively, [birds [that fly]] [swim]]

(6') Instinctively, [[the desire [to fly]] [appeals [to children]]]

The structural descriptions of (5') and (6') reveal clearly the difference between linear and structural proximity. In both cases, "fly" is the closest verb to "instinctively" in linear distance, but the more remote in structural distance.

Examples (3) and (4) are ambiguous ("fly instinctively", "instinctively swim/appeal"), but in (5) and (6) the adverb is construed only with the remote verb, raising immediate questions: why does the ambiguity disappear, and more puzzling, why is it resolved in terms of the computationally complex operation of locating the *structurally* closest verb rather than the much simpler operation of locating the *linearly* closest verb?[11]

The principle of structure-dependence applies to all relevant constructions in all languages, as far as is known. There is a simple explanation, the only one known: linear order is not available to the internal computational system that yields syntactic structures and their semantic interpretation. If so, then linear order is a peripheral part of language, presumably introduced in externalization to satisfy conditions imposed by the sensorimotor modality that is employed (and, in fact, sign, with different sensorimotor options, uses somewhat different arrangements than speech). These sensorimotor properties may be largely or completely independent of language, thus telling us little or nothing about language.

Proceeding, the next question is why language should lack linear order, except peripherally as a reflex of the sensorimotor interface. There is a simple and plausible assumption that yields this consequence: language design is optimal; its operations follow principles of *Minimal Computation (MC).*

Specifically, the computational system of language is based on the simplest computational operation O for a recursive system: given objects X and Y already constructed, form Z = O(X, Y) without modifying X and Y or imposing any new structure in Z. In short, O is simply set-formation. In recent literature, O is called "Merge." The expressions constructed by Merge therefore lack order, and order will not be available for operations on Merge-created structures.

Expression (3) can be constructed in two different ways by iterated Merge, in one case merging *fly* and *instinctively* and in the other case merging *swim* and *instinctively* before these constructed elements are merged into the larger expression. Hence the ambiguity. The same is true of (4). In the case of (5) and (6), however, the construal rule that associates the initial adverb with the verb, again adhering to MC, will seek the closest verb, where distance is structural, linear order being unavailable. That yields the unambiguous interpretations of (5) and (6). The rules of

[11] Quite commonly, linear and structural distance coincide. That would follow for "head-first" languages like English if the process of linearization is determined by Richard Kayne's Linear Correspondence Axiom (LCA), which linearizes in terms of hierarchy. Kayne explores the matter far beyond, but we can keep to this case here.

EFTA01203408

externalization happen to place the verb in the more remote position, for reasons that apply quite independently.

Notice that the argument is the same for the standard cases in the literature on structure-dependence: auxiliary-raising, as in (7) but not (8), where *t* (*trace* in earlier literature) marks the position where the auxiliary is understood:

(7) Will birds that fly *t* swim
(8) *Will birds that *t* fly swim

Under MC, (7) is the only possibility if linear order is unavailable, while (8) would be selected if both linear order and hierarchical structure were available. The thought that (8) would express if linear order were available requires a paraphrase in language. Essentially the same argument holds for all cases of structure-dependence, in a wide variety of constructions in all languages.

The same optimal assumptions about the architecture of language yield a variety of other conclusions, some quite straightforward, some more interesting. One straightforward conclusion is that assignment of semantic roles should be order-independent; for example, the verb-object relation should receive the same interpretation in a head-initial language SVO or a head-final language SOV. That too appears to be the case over a broad range.

More interesting conclusions follow if we pursue the same reasoning further. Consider the sentences (9)-(10):

(9) [The boys expect the girls to like each other]

(10) which girls do [the boys expect to like each other]

In (9), the anaphor *each other* selects the local antecedent *the girl*, as expected under MC. In (10), however, it does not select the local antecedent *the boys*, within the bracket that is analogous to (9), but rather the remote antecedent *which girls*.[12] If we continue to assume optimal design under MC, hence that grammatical operations observe locality (minimal distance), then it follows that *which girls* is in fact the local antecedent for the anaphor. Accordingly, though what reaches the sensorimotor system is (9), the syntactic object that reaches the mind is something like (11)):

(11) Which girls do [the boys expect which girls to like each other]

Here the bracketed element is identical with (9) except that *which girls* replaces *the girls*.

The question is why language is designed in this way.

Once again, the answer is provided by the assumption that the computational rules are optimal, based on Merge. By simple logic, there are two possible cases of Merge, which we can describe

[12] This is one of the many kinds of examples that refute the proposal of Chater and Christiansen (2010) that anaphoric relations are simply ‘‘an instance of a general cognitive tendency to resolve ambiguities rapidly in linguistic and perceptual input,’’ hence do not involve language-specific properties derived from UG. This is another of the very few attempts to deal with some non-trivial property of language in such terms. It should be noted that there is valuable work integrating UG and general learning mechanisms. For example, Yang (2002).

EFTA01203409

as follows. Assume a workspace containing objects already constructed (including the minimal "atoms" of the lexicon). Select X from the workspace, then select Y to Merge to X, where Y has already been constructed. Y can either be in the workspace, external to X, or it can be a part of X (technically, a *term* of X) – external Merge (EM) and internal Merge (IM), respectively.

Sentence (9) is formed by repeated EM, yielding the appropriate hierarchical structure. To form (11), first apply repeated EM to form (9') = (9) with *the girls* replaced by *which girls.* Next apply IM merging *which girls* with (9'), yielding (11) with the appropriate hierarchical structures and with the two copies of *which girls* that yield the correct semantic interpretation.[13]

Note that there are no such notions as *Re-merge* or *Copy*; just Merge in the simplest form.

Another principle of MC yields (10) for externalization: pronounce as little as possible. At least one copy must be pronounced or there is no indication that the operations took place. Looking further, we find that either the structurally highest or lowest is chosen, depending on the construction and the language, but not other copies, for reasons that have a simple explanation.[14]

The property of displacement with deletion (Move) is ubiquitous in language, and was long considered to be a curious imperfection. That was an error (mine in particular). On the contrary, we can now see that it would be an imperfection of language if IM were not available. An approach to the phenomenon that bars IM has a double burden of justification: it must justify the stipulation barring IM and must also justify whatever new mechanisms are designed to yield what comes free under IM, assuming MC. The "copy theory of movement" illustrated in (9)-(11) yields quite intricate semantic interpretations (called "reconstruction" in earlier work).

Throughout, the results follow from the assumption that the design of language keeps to the overriding conditions MC. For these cases at least, UG reduces to providing a combinatorial operation to permit recursive generation of structures that provide semantic-pragmatic interpretations (and secondarily, can be externalized).

The construal of the anaphor, as in (9)-(11), keeps to minimal structural rather than minimal linear distance, as illustrated in (12), again suggesting that linear order is not available for the internal computational system:

(12) Women with children like each other

Further inquiry into anaphoric relations yields many intricacies, discussed in a rich and expanding literature, but elementary properties such as these appear to hold quite generally, in one or another form.

I mentioned that the two types of modules – input, central -- suggest two different ways of approaching the nature of language: as a parsing system (with production and the linkage to input left to the side) or as an internal cognitive system accessed by various uses of language. The considerations just reviewed bear directly on this question.

[13] I ignore here the insertion of *do*.

[14] For discussion, see Chomsky (2013a).

EFTA01203410

Let's continue to keep to the assumption that whatever the computational system of language is, it keeps to the overriding principle MC as far as possible. That makes good sense on general grounds of scientific method and also with regard to origin of language, a guiding concern since the early days of generative grammar, contrary to much misunderstanding.[15]

Suppose that language is fundamentally an internal generative cognitive system accessed by various uses of language. UG determines that language incorporates a combinatorial operation, and by MC, it is the simplest one possible (Merge). We then have an explanation for the properties of language of illustrated above: (a) the ubiquitous property of displacement, along with important steps towards semantic interpretation of constructions with displacement; (b) apparent violation of locality with anaphora; (c) the universal property of structure-dependence of rules. Further recourse to the overriding principle of MC determines that what reaches the ear has gaps that have to be filled by the parser – in the case of (10), the missing phrase (*which girls*) that receives the same semantic role as the overt phrase *the girls* in (9) as subject of *like each other*, and by the same mechanism, serves as the local antecedent for *each other*. In this case the parsing problem is fairly simple, but locating the gap and filling it ("filler gap problems") can be quite complex because of the deletion of the copies mandated by MC.

Suppose, in contrast, that language is fundamentally a parsing system. Then all of these properties remain a mystery. In the many varied cases of structure-dependence, for example, we would expect that parsing would make use of the simple computational procedure of minimal linear distance rather than the complex procedure of minimal structural distance, contrary to fact in all relevant cases in all languages. Similar observations hold for the other cases discussed.

Note again that language design seems to pose numerous problems for parsing, in particular, the familiar filler-gap problems illustrated in a simple form in (10). The same conclusion is supported by numerous other familiar cases: structural ambiguity, garden path sentences, many island properties. These seem to arise by allowing rules to run freely, posing problems for parsing – and hence also for communication, which, for many reasons including these, does not appear to have the central role assigned to it in much modern doctrine. In fact, in all cases that I know of where communicative and computational efficiency conflict, the latter is selected, as in the examples illustrated above.

The evidence, then, strongly suggests that language is fundamentally an internal generative module providing the means for construction and expression of thought, with ancillary operations of externalization reflecting properties of the sensorimotor system, pretty much along traditional lines.

Fodor cites Hilary Putnam's 1961 suggestion (p. 50) that "there are grammatical transformations because communicative efficiency is served by the deletion of redundant portions of messages, etc.)." At the time, there were, understandably, many such suggestions about why language should have the odd property of displacement (hence grammatical transformations, or some other mechanism to deal with the "imperfection"). The situation has been different for some years, ever since it has been understood that displacement and its analysis in terms of IM is to be expected on the simplest assumptions, and that problems would arise if languages lacked this property. We can rephrase Putnam's suggestion in current terms as the thesis that deletion rules

[15] Cf. Lenneberg, *op. cit*. For some discussion, see Chomsky (2014), Berwick and Chomsky (2016).

EFTA01203411

apply in externalization to enhance communicative efficiency. Insofar as that it is plausible, they enhance the efficiency of *production*, but at the same time cause difficulties for parsing by posing filler-gap problems, as in the case of the obligatory deletion illustrated in (10). The suggestion again suggests that parsing is a peripheral aspect of language.

It is important to recognize that there is compelling evidence from neuroscience and psycholinguistics supporting the conclusion that linear order is not available for the computational system.[16] It should be clear that if these conclusions about the general architecture of I-language are generally accurate, then a good deal of the technical work in linguistics must be reconsidered,[17] along with much general thinking about language and its functions and evolution.

These conclusions, if correct, imply nothing about the significance of the modular approach to parsing, and input operations generally, that Fodor develops. Rather, they place the study of parsing within the general domain of perception, with application to language a special case that may not be particularly informative about the nature of language.

Parsing is a form of behavior, and accordingly involves many different factors, of which the role of the language is only one. Hence the study of parsing seeks to identify the contribution of the language of the person carrying out this activity and to extricate it from the complex. As Fodor puts the point (135), "something like a representation of a grammar for L must be contained" within the parser, even for assigning tokens to types, surely beyond.[18] And the same L must be contained within the production system. That raises the question of what L is, if it is not an I-language in the sense of the acquisition-based approach to modularity. The latter approach focuses directly on the person's I-language, and is free to use all sorts of evidence to determine what this system is, without limit, as in the sciences generally. But if L is not a central module of the kind discussed here, questions arise about what it is, how we discover its properties, and how it fits into the general cognitive architecture.

The inquiry into parsing requires that we distinguish performance from competence; we distinguish actual behavior from generation by the linguistic system "contained" within the parser, in Fodor's terms. This distinction is often regarded as contentious, though it should not. Whatever organic system we are investigating, we want to determine its intrinsic nature and how this enters into its various uses – in this case, to determine how a person's I-language enters into parsing and other uses of language.

The distinction, which is implicit in traditional grammar, came to the fore as soon as the earliest efforts were undertaken to construct generative grammars. A familiar example is embedding, in the interesting case, with nested dependencies. As observed 50 years ago, without external aids (time, paper, pencil, etc.), sentences can be recognized as clearly grammatical with about six nested dependencies, while disruption of one of the dependencies (say, by replacing an occurrence of "if" by "either") renders it ungrammatical. With external aids there is of course no

[16] Musso et al., (2003), following the paradigm of Smith and Tsimpli (1995). For replications, see Moro (2013). Smith (2004). Costa and Lobo (2015).

[17] And, correspondingly, pursuit of these conclusions must deal with a great deal of linguistic work that appears to be inconsistent with them.

[18] See Chomsky (1965), I.2.

EFTA01203412

relevant bound on nesting.[19] Linguistic competence is not bounded by memory, though performance, such as parsing, of course must be. There is a simple explanation for the fact that parsing decays with increased nesting, and reaches a limit (without external aids) at about 7: Miller's famous "magic number" (Miller 1956). Actual speech naturally tends towards parataxis, so embedding rarely goes beyond 2. Hence the I-language property of unbounded nesting (like core properties of language generally) is not acquired by some kind of data-processing but rather derives from inherent properties of the language faculty, from UG, a part of Hume's hidden hand of Nature that enters into all forms of learning and growth.[20]

The situation is similar to arithmetical competence, which for some reason is considered less contentious. No one is confused about the fact that we can only add small numbers "in our heads," but can go on indefinitely with external aids. In brief, both language and arithmetic are based on the Turing architecture that Fodor describes, part of their essential nature, possibly with common roots (see Chomsky 2010).

It remains to consider the apparent contradiction between the postulation of central modules and Fodor's thesis that central systems lack any modular structure, but rather are "Quinean and isotropic." The contradiction is only apparent. Fodor is concerned with the central processes of fixation of belief, which indeed have the properties he describes. But knowledge of language (linguistic competence, having an I-language) is not some kind of structure of beliefs.

My one real disagreement with Fodor's account is his opening section (3ff.) on what he calls "neocartesianism," "what [Chomsky] means": namely, that the I-language a person acquires is "a body of innate *propositional attitudes*" (129). But I have never meant anything of the sort, and agree with Fodor that the idea makes little sense.[21] A person whose I-language has the properties discussed above may have all kinds of beliefs about the expressions used as illustration here, or about his or her language. Some might be true, some false, but they are not what constitutes the language that the person has mastered and uses, any more than in the case of the visual system or insect navigation.

The confusion pretty clearly arises from Fodor's interpretation of the phrase "knowing a language," the normal locution in English (not other languages) for what in more technical terms we might call having internalized an I-language. As Fodor remarks, "knowledge is – or so many philosophers tell us – inter alia a normative notion having much to do with standards of justification." It is true that "so many philosophers tell us" in discussion of propositional knowledge, but the comment clearly does not hold of normal English usage, including the case in question. When one says, for example, "I know many of my cousins, I know their flaws and

---

[19] Miller and Chomsky (1963), Chomsky (1965). Self-embedding has much narrower restrictions. See these sources and Chomsky (1964) for some early proposals on a parsing principle which, it seemed, might also account for what was later called "the wh-island constraint."

[20] There is a great deal of rather surprising confusion about these matters in current technical literature. See the introduction to the 2015 reprinting of Chomsky (1965) for some discussion. See Chomsky (2015a) on fallacious arguments in the technical literature seeking to refute the trivially obvious observation of Chomsky (1956) that unbounded nested dependencies cannot be accommodated by finite automata, the standard models of the time.

[21] Separately, I think Fodor's interpretation of "Cartesianism" in similar terms is dubious. The Cartesian concept of innate ideas, discussed briefly above, does not seem to be properly interpreted in terms of propositional attitudes. I think we can also question Fodor's interpretation of Hume's "epistemic boundedness" as relying on his "Empiricist theory of *meaning*" (124). Hume's conclusion that "mysteries of nature" lie "in that obscurity, in which they ever did and ever will remain" has quite distinct sources. See Chomsky (2009, 2013c).

EFTA01203413

foibles, I know some of the reasons for them, I partially know their languages but I don't know the rules of verbal morphology though of course I know the rule of structure-dependence," etc., there is no reason to seek a tortured, irrelevant, and hopeless account in terms of knowing-that or knowing-how, of propositional content, networks of belief, etc. That's not what the phrase "knowing X" means. And invoking subdoxastic beliefs (whatever their merit in other contexts) does nothing here but deepen the confusion.

Fodor observes that "Chomsky himself is quite prepared to give up the claim that the universal linguistic principles [say, structure-dependence] are innately *known* in favor of the explicitly neologistic (hence sanitized) claim that they are innately `cognized'," but he misconstrues the reasons. It is simply an effort to avoid pointless debates with philosophers who insist on taking refuge in Wittgenstein's fly-bottle instead of using the terms of ordinary language with their own meanings (as in "knowing X"), or, as is commonly done even in the hard sciences, using these terms intelligibly if sometimes laxly in informal discourse. My point was much the same as Turing's in his famous paper introducing the imitation game, where he warned that the question whether machines can think "is too meaningless to deserve discussion" – along with such questions as whether kites fly, submarines swim, Chinese rooms translate, etc. These are questions of ordinary usage in one or another language, or, sometimes, of what metaphorical extensions we choose to make. They are not substantive questions.

When we put these misinterpretations aside, there is no contradiction between the postulation of acquisition-based central modules and Fodor's rejection of central modules for fixation of belief.

In his discussion of the "Quinean" and "isotropic" character of internal systems -- meaning that any evidence is in principle relevant -- Fodor states that "some linguists" deny this property for language, claiming that "no data except certain kinds of facts about the intuitions of native speakers could, in principle, be relevant to the (dis)confirmation of grammatical theories." If so, they are severely mistaken. The only advocacy of this restriction that I know of is by Quine, who repeatedly insists that "there is nothing in linguistic meaning" (which he construes to extend to properties of syntax and semantics generally) "beyond what is to be gleaned from overt behavior in observable circumstances," proceeding to restrict the latter to "querying sentences for assent and dissent."[22] Quine's restrictive stipulations contrast sharply with the practice and principles of generative grammar from its modern origins, which always insisted that evidence of any kind is in principle relevant to "the (dis)confirmation of grammatical theory," including evidence from other languages, available once we recognize the role of the species property UG.

Quine's restriction of relevant evidence is part of a much broader thesis, which might merit a few words in the light of its great influence and what it tells us about the tenor of the times. Quine's guiding principle in this domain is summarized clearly in his *Pursuit of Truth*: "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" (Quine 1990, 37). An analogous argument would be that the study of the visual system must restrict itself to the visual stimuli that determine the specific form that the visual system assumes. Of course, that argument would be dismissed at once: though indeed the mature visual system is a function of input stimuli (and as well known, it can vary substantially

[22] Quine (1975, 1992, p. 46)

EFTA01203414

depending on stimulation in early infancy), the outcome depends on many factors, including genetic endowment, and the scientist studying the visual system is free to consider these and indeed whatever evidence might be relevant to how the organism grows. But these options are barred in principle to the linguist, on the tacit assumption that the language faculty cannot in principle have any basis in human biology – that there can be nothing like UG. The linguist cannot in principle then learn anything about English from the study of Chinese, or from psycholinguistics, or neuroscience, or any other source. The central system of language (if that's what it is – what else could it be?) violates the Quinean and isotropic properties of central systems.

Note that (dis)confirmation of a *theory of language* (*or of particular I-languages*), relying on any evidence in principle, is not to be confused with the operations of language acquisition. In this case, to quote Fodor again, "the neural mechanisms…develop according to specific, endogenously determined patterns under the impact of environmental releasers"; and as in the case of growth and development of other subsystems of the organism, only certain "environmental releasers" trigger and shape the process.

To summarize briefly, I think Fodor is right to recognize two mental processes that are "basically a reflex": the input modules that are his topic and acquisition of language (with the qualification mentioned earlier), the latter providing a central module that falls together with others, but is not a system of propositional attitudes acquired by fixation of belief. This central module is accessed for production (occasionally externalized) and parsing. The latter, like all of performance, is a mixed system guided in some manner by the internal language but involving many other factors. The central module itself is a biological object, whose nature we seek to discover, using any evidence available, with no such restrictions as those that Quine imposes. The two approaches suggest two ways of seeking the fundamental nature of language. There is substantial evidence, I think, favoring the latter, which has something of a traditional flavor. If the approach outlined here is on the right track, then considerable rethinking of the nature and use of language is in order, both within technical linguistics and in reflections on its nature and use.

Biblio

Albright, T. (2015). "Perceiving," Daedalus, Winter 144(1).

Berwick, R., and N. Chomsky (2016). Essays on the Evolution of Language. MIT.

Bizzi, E., and R. Ajemian. (2015). "A Hard Scientific Quest: Understanding Voluntary Movements." Daedalus, Winter 144(1).

Chater, N. and M. Christiansen (2010). Language Acquisition Meets Language Evolution. Cognitive Science 34, 1131--1157.

EFTA01203415

Chomsky, N. (1956). Three Models for the Description of Language. I.R.E. Transactions on Information Theory IT-2, 113-124.

Chomsky, N. (1964). Current Issues in Linguistic Theory. Mouton.

Chomsky, N. (1965). Aspects of the Theory of Syntax. MIT. Reprinted with new introduction, MIT, 2015.

Chomsky, N. (1975). Reflections on Language. Pantheon.

Chomsky, N. (1986). Knowledge of Language. Praeger.

Chomsky, N. (2009). The Mysteries of Nature: How Deeply Hidden? J. of Philosophy, April. Reprinted in Chomsky (2015b).

Chomsky, N. (2010). Some simple evo devo theses: how true might they be for language? In R. Larson, V. Deprez, and H. Yamakido, eds., The Evolution of Human Language: Biolinguistic Perspectives. Cambridge.

Chomsky, N. (2013a). Problems of Projection. Lingua 130.

Chomsky, N. (2013b). What is Language? J. of Philosophy, December. Reprinted in Chomsky (2015b).

Chomsky, N. (2013c). What can we Understand? J. of Philosophy, December. Reprinted in Chomsky (2015b).

Chomsky, N. (2015a). A Discussion with Naoki Fukui and Mihoko Zushi. In Sophia Linguistica 64. The Sophia Linguistic Institute for International Communication (SOLIFIC), Sophia University, Tokyo.

Chomsky, N. (2015b). What Kind of Creatures are We? Columbia.

Chomsky, N. (2016). Language Architecture and its Import for evolution. To appear in Berwick and Chomsky (2016).

Churchland, P. (2013). Introduction, W.V.O. Quine, Word and Object. MIT.

Costa, J. and M. Lobo (2015). Testing Relativized Minimality in intervention effects: the comprehension of relative clauses with complex DPs in European Portuguese. ms, FCSH/Universidade Nova de Lisboa [presented at Romance Turn, University of Islas Baleares].

Curtis, S. (2013). Revisiting Modularity: Using Language as a Window to the Mind. In M. Piattelli-Palmerini and R. Berwick, eds, Rich Languages from Poor Inputs. Oxford.

EFTA01203416

Gallistel, C. R. (1999a). "Neurons and Memory." In M. S. Gazzaniga, Conversations in the Cognitive Neurosciences, MIT.

Gallistel, C. R. (1999b). "The Replacement of General-Purpose Learning Models with Adaptively Specialized Learning Modules," in M. S. Gazzaniga, ed., The Cognitive Neurosciences, 2nd edition, MIT.

Hoffman, D. (1998). Visual Intelligence. Norton

Lenneberg, E. (1967). Biological Foundations of Language. Wiley.

Miller, G. A. (1956). The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. Psychological Review 63.81-97.

Miller, G. A., and N. Chomsky (1963). In R Duncan Luce, R. W. Bush, & E. Galanter, eds., Handbook of Mathematical Psychology II.269-322, Wiley.

Moro, A. (2013). The Equilibrium of Human Syntax. Symmetries in the brain, Leading Linguists Series. Routledge.

Musso, M., A. Moro, V. Gluache. M. Rijntjes, J. C. Bücheli, and C. Weiller (2003). Nature Neuroscience 6.774–781.

Quine, W. V. O. (1975). "Mind and Verbal Dispositions." In S. Guttenplan (ed.), Mind and Language. Oxford.

Quine, W. V. O. (1990). Pursuit of Truth. Revised edition, 1992. Harvard.

Smith, N. (2004). Chomsky: Ideas and Ideals. Cambridge.

Smith, N., and I. Tsimpli (1995). The Mind of a Savant: Language Learning and Modularity. Oxford.

Tomasello, M. (2009). Universal Grammar is Dead. Behavioral and Brain Sciences 32:5.

EFTA01203417