Faculty of Language, Functional Models, and Mechanisms*,**

Liza Skidelsky

Universidad de Buenos Aires;
National Scientific and Technical Research Council (CONICET)
lskidelsky@filo.uba.ar

In a recent series of papers, John Collins has challenged the dominant epistemic view of Chomsky's faculty of language (FL), which holds that the FL consists fundamentally of propositional knowledge. Collins presents the architectural view that holds that the FL is a computational information-processing system. I fully endorse this broad architectural perspective. Nonetheless, I would like to discuss one aspect of his architectural view which maintains that we should not understand the FL as a causal mechanism, that is, as part of a causal nexus. In this paper, I will try to develop the main lines of an alternative, though perhaps broadly compatible, way of unfolding the architectural perspective in which it makes sense to think of an aspect of the minimalist program as a cognitive functional model that nevertheless describes a causal

mechanism. I will argue that there are no compelling reasons to discard the possibility of conceiving the FL as a causal mechanism (albeit an idealized one) of the same nature as the mechanisms which any scientific theory about cognitive architecture attempts to explain. The model in and of itself is not a mechanistic one in that it only specifies the functional properties of its object of description, leaving aside structural properties such as location, temporal order of processing, and the like. Still, the object being described, the FL, can be conceived of as a mechanism. Unlike the advocates of mechanistic explanations, I will argue that there are cognitive mechanisms that can have a genuine functional explanation (i.e. that do not constitute a 'mechanism sketch') depending on the correspondence that can be achieved between the cognitive model and the cognitive mechanism. If the correspondence between the entities, activities, and organization postulated by the cognitive model is direct or straightforward regarding the entities, activities, and organization of the mechanism, then there are good chances of obtaining a mechanistic explanation; one in which not only the functional, but also the structural, properties of the mechanism are specified. If, on the other hand, the correspondence is indirect, as in the case of the FL mechanism, the functional explanation appears to be the most adequate to the extent that it highlights the relevant explanatory characteristics of the mechanism.

   Key words: *language faculty, mechanism, competence-performance distinction, functional explanation*

1. Introduction

The epistemic thesis of the nature of the faculty of language (FL) is by far the most widespread conception in the field of philosophy of language and philosophy of linguistics. The epistemic thesis understands the FL to be the propositional knowledge that a speaker/hearer possesses about her language (Fodor, 1983; Dwyer and Pietroski, 1996; Knowles, 2000). To counter this view, Collins (2004, 2006, 2007, 2008) proposes an architectural thesis which considers the FL to be a computational information-processing system. I fully agree with Collins in his rejection of the epistemic conception. I think Collins did a wonderful job making clear, mainly, why thinking of the FL as a propositional database results in a loss of explanatory power of generative linguistics, so I will not address this matter here, but rather briefly in

section 2.1. As a consequence, I also endorse his architectural proposal in a broad sense. Nevertheless, I would like to address one aspect of his architectural view which holds that we should not understand the FL as a causal mechanism. His main reason is that since the FL is characterized in terms of "an abstractly specified computational (= intensional function) system of the mind/brain[…] This characterization renders the faculty neither as a set of propositions to be known, nor as a mechanism, a part of a causal nexus" (Collins 2004, 529-530).

I will try to develop the main lines of an alternative, although perhaps broadly compatible (we will see in what sense and extent), way of unfolding the architectural perspective in which it makes sense to think of an aspect of the minimalist program as a cognitive abstract functional model that nevertheless describes a causal mechanism. I will argue that there are no compelling reasons to discard the possibility of conceiving the FL as a causal mechanism (albeit an idealized one) of the same nature as the mechanisms which any scientific theory about cognitive architecture attempts to explain. The FL model in and of itself is not a mechanistic one in that it only specifies the functional properties of its object of description, leaving aside structural properties such as location, temporal order of processing, and the like; in this sense, I completely agree with Collins that the theory of the language faculty describes it as an 'abstractly specified' structure of the mind/brain. Still, the object being described, the FL, can be conceived of as a mechanism; this would be my main point of disagreement with Collins, although it depends on how this question is spelled out.

I will therefore defend an architectural view according to which cognitive psychology's models, which I believe are mostly functional models, aim to account for cognitive mechanisms. Accordingly, I would like to stress a point that, although it may be a truism, it is often overlooked in philosophy of cognitive science and that also seems to appear in the debate about the FL. At the bottom of the viability of offering an alternative architectural view is the dismantling of the inference which spans from the contention that the theory of the FL is about abstract aspects of the mind/brain to the affirmation that it does not describe a causal mechanism. This inference seems to be a case of the inference held by most advocates of functional explanation in cognitive science against mechanistic explanation. The

defenders of functional explanation consider that cognitive explanation does not capture mechanisms (Fodor 1968, Cummins 1983, Weiskopf 2011). Their argument is based on the fact that since psychological explanation, which makes use of cognitive models, is not a mechanistic but rather a functional explanation, it does not account for mechanisms. In other words, since cognitive theories or models only describe functional properties of the entity responsible for the cognitive phenomena to be explained, that is, "activities or manifestations of their causal powers, dispositions, or capacities" (Piccinini & Craver 2011, 291), leaving aside the structural properties such as localization, size, orientation, spatio-temporal organization, etc., this entity cannot be a causal mechanism.

This inference is based on what I call an ontic-epistemic assertion, which is often at work in the debate between functional and mechanistic explanation. The assertion consists of considering that the kind of explanation one should offer depends on the kind of entity responsible for the cognitive phenomena to be explained; for example, if what there is in the world is a mechanism, the adequate explanation should be mechanistic. Or, to frame this assertion conversely, the kind of entity responsible for the cognitive phenomena to be explained is determined by the kind of explanation one offers; if the explanation is functional, the entity cannot be a mechanism. Accordingly, defenders of both mechanistic and functional explanations link mechanistic explanations with the notion of mechanism. And, while the advocates of mechanistic explanations consider that psychological explanations are a kind of mechanistic explanation since they aim to describe mechanisms (Piccinini & Craver 2011), the advocates of functional explanations consider that psychological explanations are not mechanistic because they do not capture mechanisms (Weiskopf 2011).

To argue against this ontic-epistemic assertion would involve a detailed analysis of both sides of the relation so as to show that there is no determination or relation of dependence in any direction. For the topic at hand, it is sufficient to show that there are no compelling reasons to discard the possibility that the responsible entity of cognitive phenomena is a mechanism. However, we are not therefore committed to the view that the explanation has to be mechanistic as we likewise are not required to commit ourselves to the view that cognitive explanations do not capture mechanisms in order

to defend some non-mechanistic (i.e. functional or abstract) form of psychological explanation or specification. Thus, if I can make the case to take into account a truism that is often forgotten in these debates (the distinction between the ontological level of entities in the word and the explanatory or broadly epistemic level of explanations, descriptions, interpretations, characterizations, or specifications of those entities), the fact that there are no compelling reasons to discard the possibility that the responsible entity of a cognitive phenomenon is a mechanism (in this case, the FL) does not commit us to a mechanistic explanation, model, interpretation, etc., of this entity.

In this sense, unlike the advocates of mechanistic explanations, I will argue that there are cognitive mechanisms that can have a genuine functional explanation (i.e. that do not constitute a 'mechanism sketch') depending on the correspondence that can be achieved between the cognitive model and the cognitive mechanism. If the correspondence between the entities, activities, and organization postulated by the cognitive model is direct or straightforward regarding the entities, activities, and organization of the mechanism, then there are good chances of obtaining a mechanistic explanation; one in which not only the functional, but also the structural, properties of the mechanism are specified. If, on the other hand, the correspondence is indirect, as in the case of the FL mechanism, a functional explanation appears to be the most adequate to the extent that it highlights the relevant explanatory characteristics of the mechanism.

I would like to be clear about my purpose in this paper so as to avoid misunderstandings as much as possible. I will not try to show that the FL is a mechanism. Instead, I will defend a much more modest claim, which is a modal claim that deals with possibility and is based only on the discussion of Collins's architectural proposal. My claim is that Collins's reasons for ruling out the possibility that the FL may be a mechanism are not compelling. In this sense, I will only concentrate on discussing those reasons that I think one can find in Collins's defense of the idea that the FL is not a causal mechanism. Of course, if one wished to defend the stronger thesis that in fact the FL is a mechanism, rather than questioning only these reasons or relying on the programmatic assertions made by Chomsky to the effect that it is possible to conceive of it as such, one should look at the linguistic theory itself and ask how such a possibility might go. Furthermore, one would

also have to look at how the computational mechanism would work and, in this sense, try to offer a processing construal of the FL function. But, I insist, my claim is a much weaker one; it is a modal claim about possibility and is completely based on the reasons Collins offers for believing that the FL cannot be a causal mechanism, accompanied by some general considerations in favor of such a possibility.

Nor will I try to discern a causal mechanistic–albeit very abstract–interpretation of linguistic theoretic descriptions of the FL, since I think that a functional interpretation is the most adequate (section 3.1); a functional interpretation that, in this case, is not, and probably cannot be, a kind of mechanistic explanation (section 3.2). Hence, I will also not try to show that the description provided by linguistic theory of the FL is a kind of functional description that would support a causal mechanistic interpretation. Unlike defenders of mechanistic explanation, I do not think that functional explanations are a kind of mechanistic explanations (section 3.2). Consequently, there is no need to look at the linguistic theory itself to ask how a mechanistic interpretation of the FL might go. In this sense, as I have mentioned before, I fully agree with the abstract specification that Collins offers from his architectural view about the FL. I will simply try to show that that same specification does not conclusively preclude conceiving of the entity responsible for linguistic phenomena as a mechanism. In this sense, just as Collins does not provide reasons specific to linguistic theory itself for defending his view that the FL is not a causal mechanism, I will also not provide them in order to defend the view that Collins's reasons are not conclusive to rule out this possibility. Of course, it might be interesting if there are reasons specific to the linguistic theory itself for ruling out the possibility that the FL could be a mechanism. However, I will not address this issue here to the extent that, as I have mentioned, it is not my aim to show that the FL is a mechanism, but rather to argue that the reasons offered by Collins are not conclusive for ruling out the possibility in question.

By keeping in mind the truism of the distinction between the ontological level of entities in the word and the epistemic level of explanations, descriptions, interpretations, characterizations, or specifications of those entities, we may attempt to dismantle the ontic-epistemic assertion. This will allow us to maximize the architectural view or, in other words, to propose what

can be viewed as a bold architectural view in contrast to—or perhaps, ultimately compatible with—the much more cautious and careful version of the architectural view offered by Collins.

The structure of the paper is as follows: in section 2, I will address the question of the abstract specification of the FL and the possibility of conceiving of the FL as a mechanism. First, I will present Collin's abstract specification of the FL, with which I completely agree (2.1). Second, I will discuss one particular aspect of his architectural view. I will concentrate on the reasons that I believe Collins offers for arguing that, since the specification of the FL is abstract, the FL is not a causal mechanism. My aim in this section is to show that these reasons can be compatible with the possibility that the FL is a causal mechanism (2.2). Finally, I will consider some reasons that could support the idea that it makes sense to conceive of the FL as a mechanism (2.3). I want to stress again that I will not try to prove that the FL is, in fact, a mechanism (this is something that, ultimately, I doubt anyone can do, as with any attempt to prove an ontological fact). I will only try to show that there are no compelling reasons to dismiss this possibility. In section 3, I will connect this discussion with the debate in the philosophy of cognitive science between functional and mechanistic explanation. First, I will try to specify in what sense I consider the FL model as a functional model of a mechanism (3.1), and second, I will address the reasons why the FL's explanation is not mechanistic despite the fact that the FL can be considered a mechanism (3.2).

2. The Faculty of Language as a Computational Mechanism

The most widespread conception of the nature of the FL, at least in the field of philosophy, is the epistemic thesis. This thesis, which has been maintained primarily by Fodor (1975, 1983, 2001), affirms that linguistic competence or I-language is about what the speaker/hearer *knows*. Since knowledge is traditionally understood in terms of propositional attitudes, linguistic competence would consist of a set of propositional attitudes (i.e. beliefs) about the particular language attained by the speaker/hearer.[1] In this

---

[1] This is the most representative position of the epistemic thesis. This conception

way, the thesis of the FL, according to Fodor and expressed by Collins (2004, 503), is not "a proposal about the *architecture* of the mind/brain, where such a thesis offers a causal explanation of *how* speaker/hearers acquire and maintain a knowledge of language and put it to use."

In contrast to this conception, the architectural version considers the FL to be a computational system of the mind/brain in charge of linguistic information-processing. Collins (2004, 2006, 2008, 2009) defends a particular view of this architectural thesis which maintains that although the FL is considered a computational system, to the extent that it is specified abstractly, it is not part of the causal structure of the mind/brain (Collins 2004, 530). This means that although the FL is an aspect of the mind/brain, it is abstracted from its causal structure. Thus, in favor of his abstract architectural version, Collins holds that to the extent that its description is not in causal terms, the FL is not a causal mechanism. I will first present the standard abstract description of the FL (2.1); second, I will address the reasons I think Collins offers for arguing that since the specification of the FL is abstract, the FL is not a causal mechanism (2.2); and third, I will present some reasons that could support the idea that, in light of recent developments in the minimalist program, it makes sense to conceive of the FL as a mechanism (2.3).

2.1. The Abstract Specification of the Faculty of Language

Collins (2004, 2006, 2008) is very careful and precise in presenting his version of the abstract architectural thesis following Chomsky's development of the minimalist program (Chomsky 1995a). Collins does not affirm that his abstract architectural conception is the one that, in effect, Chomsky sustains, but rather that it aligns with the Chomskyan developments and is independently coherent. However, I think he reflects the main tenets of the FL theory very accurately. This presentation makes clear that the emphasis on the characterization of the FL is not placed on the descrip-

---

covers all the positions that consider that the I-language consists (only or fundamentally) of a body of linguistic information in the mind/brain. This linguistic information can be understood from a database, on through mental representations and propositional attitudes, up to knowledge.

tion of the inventory of categories and structural relations among linguistic representations, but rather on the computational mechanism that computes this linguistic information. As is well known, the FL is a subsystem of the human mind/brain. It has two main components: a cognitive system and performance systems (articulatory-perceptual and conceptual-intentional). Linguistics is concerned with the initial state and the steady state of the cognitive system. Universal Grammar (UG) is the study of the initial state, and the 'grammars' are related to the study of I-languages, that is, the steady state achieved by the speaker/hearer, or her linguistic competence. UG consists of a finite set of principles valid for all possible human I-languages and parameters which are fixed in virtue of linguistic experience during language acquisition. Both principles and parameters are part of the human biological endowment.[2]

I-languages are different instantiations of the UG's parameters. The technical notion of I-language assumes a certain idealization that consists of speakers with uniform experience in homogeneous linguistic communities (without diversity of dialects or variation among speakers). The I-language is Internal (to the mind/brain), Individual (to particular speakers) and Intensional (a function characterized in intensional terms that generates structural descriptions). The I-language (both in its initial state and its steady state) is a real object in the minds/brains of the speakers; the linguist studies those abstract aspects of the brain, that is, the mental aspects, which in this case are the linguistic aspects. Insofar as its object of study is a natural object, linguistics is part of the natural sciences (since it is part of psychology, which is ultimately part of the field of biology) and addresses its object of study with the same methodology used by the rest of the natural sciences (this is the thesis of methodological naturalism).

According to the minimalist program, each I-language consists of a

---

[2] Principles are restrictions which every possible human language is subject to. For example, the object of a verb forms a phrase with the verb that does not include the subject, or put another way, objects combine with verbs before subjects do (verb-object requirement). Parameters, in contrast, are fixed by linguistic experience. An example is the Null Subject Parameter, whose setting determines whether a tensed clause must have an overt subject noun phrase, as in French, or need not, as in Spanish (Baker 2001).

computational mechanism and a lexicon (Chomsky 1995a, 1995b). The (intensional) function of the mechanism is to generate structural descriptions of linguistic expressions from lexical choices. The computational procedure takes lexical features as inputs, and its outputs are the structural descriptions, which are abstract (symbolic or formal) objects that carry all the semantic, syntactic, and phonetic information required to associate the sound and meaning of linguistic expressions. The lexicon is a list of idiosyncratic lexical items which are not derived from general principles. These lexical items are characterized by a set of semantic, syntactic, and phonetic features. The computational mechanism computes the phonetic features of a lexical item to bring about a phonetic representation (of the phonetic form, PF) and computes the semantic features to bring about a semantic representation (or logical form, LF). The formal features that participate in the computations should be eliminated on the way to the PF to bring about a convergent representation, that is, one interpretable by the interface performance systems.

It is these interface systems that impose conditions of legibility since they interpret the outputs and follow their instructions for thought and action. The computational mechanism computes the features of lexical items (by means of the operation *Merge*) until reaching a point denominated *Spell-Out*, in which the derivations separate toward the two interface levels. The operations that are carried out between *Spell-Out* and the PF are not of the same nature as those which bring about the LF; it is assumed that the computational mechanism leads to the LF, whereas a separate phonological component performs the operations to bring about the PF. Within the set of convergent derivations that satisfy the interface conditions, there is a subset of admissible derivations selected by the system's principles of economy. The output of the I-language for a linguistic expression E is, at least, the pair {PHON, SEM} formed by an optimal convergent derivation. PHON(E) and SEM(E) are representations that carry information about the phonetic and semantic properties of E, respectively. PHON(E) and SEM(E) undergo interpretation by the articulatory-perceptual and conceptual-intentional systems, respectively.

In this sense, the FL is a subsystem of the mind/brain whose function is to assign structural descriptions to linguistic expressions. It has an initial,

genetically-determined state, and through computations of this initial state plus the data provided by the linguistic environment, the FL brings about outputs that will be interpreted by other subsystems that underlie the use of language. Thus, it is clear from the standard abstract description of the FL that it is a computational system of linguistic information-processing. Conversely, understanding the FL only as a body of knowledge or merely as a database or a set of linguistic representations would imply the loss of the explanatory power of the model to account for sound-meaning pairings (both in language acquisition and in the mature workings of this capacity).[3]

Based on this standard general characterization of the FL, Collins (2004) defends the idea that although it is an aspect of the brain, insofar as its description is not couched in causal terms, the FL is abstracted from the causal structure of the brain, and, in that sense, we cannot consider it to be a causal mechanism. In the next subsection, I will discuss his reasons for maintaining that claim in order to show that they are not conclusive to reject the possibility of conceiving of the FL as a causal mechanism, one which is at work whenever a speaker/hearer operates with linguistic representations.

2.2. The Faculty of Language as an Abstract Computational System

I suppose that Collins would not disagree with the claim that it is possible that the theory about the FL has a role to play in a future integrated theory of linguistic performance and, moreover, that it could be integrated into cognitive neuroscience (not without qualifications: see section 3). Perhaps what is at the bottom of the abstract version of the architectural thesis is a certain caution about how the theory of the FL can fit in with the possibility that the FL is a mechanism that is a part of the cognitive causal architecture. This concern is often couched in what I advance in section 1 as an inference that begins with the affirmation that the theory of FL offers an abstract specification and concludes that its object of study, the FL, is not a causal

---

[3] Of course, there are additional reasons other than preservation of the model's explanatory power for defending that the FL does not (only or basically) consist of a body of information. I will not address this issue here. See Skidelsky (forthcoming) for a discussion of some of these reasons, and, naturally, Collins (2007, 2008).

mechanism. This inference, in the field of the philosophy of linguistics, is a case of a more general inference that is quite widespread in the philosophy of cognitive science. In particular, the defenders of the functional explanation hold that psychological explanations (or cognitive models) are not explanations (or models) of mechanisms and therefore do not describe mechanisms (Fodor 1968, Cummins 1983, Weiskopf 2011).

I believe that this direct relation between the kind of explanation (or model) and the entity postulated to be responsible for the cognitive phenomena being explained (or modeled) is misleading. There do not seem to be good reasons either for considering that an abstract description of a mechanism establishes by itself that the description is about an abstract object (in contrast to a causal mechanism) or, as we will see in section 3, for considering that the specification of a mechanism should be, by itself, mechanistic. In this sense, it is perfectly compatible to maintain that a certain theory or model does not offer a causal explanation of the entity responsible for the cognitive phenomena to be explained, but, nonetheless, that the theory describes a mechanism that possesses causal dispositions allowing it to be inserted into a causal nexus. Of course, this seems to be a truism. In fact, it is this truism that I want to emphasize here. In this sense, a device (e.g. a calculator) can be described as computing a function (e.g. the addition function), where such description is clearly abstract (non-causal), and yet the device described is a causal mechanism. Following these lines, but in a much weaker sense, I believe that the abstract specification of the FL does not preclude the possibility that the FL can be a causal mechanism.

As I have mentioned, in this subsection I will discuss the reasons I think Collins offers in support of his abstract architectural conception of the FL. In the next section, I will address the question of in what sense the FL model can be seen as a functional model and the more general debate in the philosophy of cognitive science between the defenders of functional and mechanistic explanation, showing its direct relation to this discussion in the philosophy of linguistics. In the same way that Collins does not affirm that his abstract architectural conception is the one that, in effect, Chomsky sustains, but rather that it aligns with the Chomskyan developments and is independently coherent, I hope to show that there are reasons for believing that the general lines of the causal architectural version which can be

inferred from my considerations can also be viewed as fulfilling those requirements.

According to Collins, "the language faculty is an abstractly specified computational (= intensional function) system of the mind/brain" (2004, 529). The theory of the FL offers an abstract specification of the function in intension that pairs sound with meaning. Although this function establishes limit conditions for what the brain has to do in order for us to be competent speaker/hearers, it does not describe the causal history of linguistic performance (Collins 2006). The explanation of how we achieve sound-meaning pairings is not a causal explanation *cum* neurological of, for example, how we are able to order a coffee, nor is it predictive of this behavior. The states of the FL are not prone to causal generalizations, and they do not contribute to the etiology of linguistic acts; hence, they do not take part in the causal connections that bring about linguistic behavior.

The FL, according to Collins, consists of a series of *sui generis* conditions satisfied by the operations of a normal brain. They are *sui generis* in the sense that they specify an aspect of a system that would not otherwise be visible. This aspect is the recursive structured integration of sound and meaning. In this sense, the faculty accounts for the competence, not the performance (a distinction introduced by Chomsky 1965). The characterization of the faculty is oriented toward the fact that having linguistic competence consists of being able to pair sounds with meanings in a systematic way. "Thus, rather than explaining performance, acts of speech or thought, the faculty hypothesis explains how performance is so much as possible" (Collins 2004, 509). Therefore, the FL would be explanatorily (not causally) related to performance in the same way that, beyond the differences, the knowledge of a musical score would explain someone's competence in regards to a particular Beethoven piece, knowledge that remains even if the person is injured and cannot play an instrument (Collins 2004).

It seems, according to these considerations, that the distinction between abstract system and causal mechanism would lie in at least two matters: the possibility of causal generalizations about linguistic behavior and the competence-performance distinction.

First we will consider causal generalizations. Collins argues that:

> The states of the faculty are solutions to an equation between the features of the lexical items and the conditions imposed by the external systems as to which features are legible to them. The states are not amenable to causal generalizations; nor, perforce, do they contribute to the aetiology of linguistic acts. (2004, 517)

One can understand this stance in the context of a discussion against a specific way of putting forward how the linguistic causal mechanism would look in order to explain linguistic behavior. Collins is arguing against Evans's (1981) and Davies's (1986) causal construal of the epistemic perspective of tacit knowledge of language. In this construal, in very approximate terms, there is a direct mapping or correspondence between the inferential roles of propositional states that carry information about a grammar and the causal roles of non-propositional states that enter into the explanation of the speaker/hearer's judgments mandated by the theory. For Collins, this construal does not work, mainly because there is not representation/represented distinction, in the sense that the FL is not represented by the mind/brain. I will not address Collins's arguments in favor of collapsing this distinction, with which I completely agree (see Collins 2007, 2008).

Instead, I want to emphasize the point that if one distinguishes the idea of a causal mechanism from the *mirror constraint*, the requirement of direct relations between the linguistic high level and the lower no-linguistic level, the possibility of conceiving of the FL as a causal mechanism still remains. I believe that the problem of this construal, beyond the one pointed out by Collins, is not so much the idea of offering a causal construal–as Collins notes–as the idea of direct relations between the mental-linguistic higher-level specifications and the lower-level specifications of a cognitive mechanism. As we will see in section 3, although the FL can be conceived of as a cognitive causal mechanism, it is highly probable that the components, activities, and organization that appear in the FL model do not have a direct relation of correspondence to the components, activities, and organization of the physical system in which it is realized.

Thus, in this context, it seems appropriate to argue against those who maintain that if the FL is a mechanism, it should explain behavior. If the reasons why the FL does not explain behavior are correct, then this specific

construal of the causal architectural approach cannot be correct. Thus, we should not understand Collins as asking that in order to consider the FL as a causal mechanism, the theory of the FL has to offer causal linguistic generalizations that cover linguistic behavior. He is simply stating that "with language, we just don't know what the generalizations would look like" (Collins 2007, 476). The problem of linguistic generalizations lies in a deep, mysterious aspect related to our cognitive incapacity of capturing why the linguistic system behaves in certain ways, giving certain answers to certain external linguistic stimuli (Collins 2006). More precisely, Collins's skepticism is reflected in this passage: "you can give an account as detailed as you like of the internal structure that would support the kind of behavior one is interested in, but one is no closer to knowing why the system behaves in one way as opposed to another" (Collins 2007, 476).

Unlike those who consider that if the LF is a mechanism, it should explain linguistic behavior, I believe that one can maintain the idea that there are not (and probably never will be) linguistic behavioral generalizations, but this does not preclude the possibility that the FL is a causal mechanism. Although the reasons why there will probably never be linguistic behavioral generalizations are different from the ones Collins offers, to the extent that he conceives the FL as a subpersonal system, perhaps he would not disagree with what I will say about the aims and scope of a subpersonal mechanism like the FL. The theory of FL specifies the function that maps a set of lexical features into pairings of sound-meaning structures that are the inputs for the performance systems. The theory in itself is not intended to account for linguistic behavior, that is, actual linguistic acts. Leaving aside the mysterious aspect that does not seem to fit well with a methodological naturalism, Chomskyan linguistic theory (or, to cite another example commonly mentioned by Chomsky himself, Marr's theory of vision) does not intend to account for linguistic (or visual) behavior, much less for the even more complex behavior in which the linguistic (or visual) capacities are involved. Linguistic behavior seems to depend on considerations that surpass the scope of computational psychology. There is an aspect related to the point of view of the agents, their interests, and other factors that intervenes in linguistic behavior and seems to go beyond an inquiry focused on cognitive mechanisms.

This is because the object of study of these inquiries is the cognitive (sub) system, the outputs of which have to come together in a way that brings about thought and action. Accounting for linguistic thought and action is a highly complex task for which there are currently neither general principles nor plausible generalizations, much less what philosophers expect, namely, *ceteris paribus* laws. The explanation of complex human behavior may be a *desideratum* to take into account in computational psychology, but its fulfillment goes beyond this science. A reasonable aim of psychology, which is observable in psychological practice, is to account for how cognitive capacities work, i.e. to account for subpersonal mechanisms (Skidelsky 2006). The fact that there are no (or perhaps never will be any) causal linguistic generalizations that cover linguistic behavior does not by itself establish that the linguistic representations–in Collins's deflationary terms (Collins 2007)–do not have causal roles in the computation machinery of the I-language. The computations described in this theory are required whenever the speaker/hearer operates with linguistic representations (Marantz 2005). Therefore, if we focus our inquiry on the mechanism of the I-language, of course there are generalizations concerning the causal dispositions that describe how this mechanism works. Moreover, the well-established linguistic theories that are concerned with phenomena that go beyond subpersonal mechanisms, such as some processing theories, do not offer such linguistic behavioral generalizations (this is the case of, for example, Garrett's model of sentence production. See Garrett 1982).

It seems, then, that the reasons in favor of the abstract architectural perspective are mainly based on the competence-performance distinction. Collins (2008) maintains that the FL theory specifies the intensional function that the underlying neuronal systems of language production and comprehension respect. This function codifies pairs of semantic and phonological representations that are explanatory, *inter alia*, of the speaker/hearers' judgments. But this conception, according to Collins, "stands in contrast to a 'processing' model of the language faculty insofar as specifying the function is not to specify *how* speaker/hearers process linguistic material" (2008, 5). Performance systems related to the production and understanding of language interface with the FL, but the FL is independently constituted; its explanation is not reducible to those systems. In fact, as Collins affirms, "it

looks as if the language faculty is not a performance system. It is that third component that is required to effect a convergence of sound and meaning, for properties of neither are predictable or explainable from the other" (2004, 518).

I would like to stress two related points. First, the competence-performance distinction is a methodological distinction, not a metaphysical one. There are no two distinct FLs. Second, Chomsky's insistence that his object of study is linguistic competence (the I-language) is not due to the fact that he does not believe the I-language mechanism should take part in a performance theory, but rather because he is of the opinion that it is the only aspect of a future integrated linguistic theory that can be studied in naturalist terms. We will first turn our attention to the second point.

Some of Chomsky's typical statements through the various developments of his theory which assert that the FL is a component of the mechanisms underlying linguistic performance are the following:

> In general, it seems that the study of performance models incorporating generative grammars may be a fruitful study; furthermore, it is difficult to imagine any other basis on which a theory of performance might develop. (1965, 15)
> If we accept–as I do–Lenneberg's contention that the rules of grammar enter into the processing mechanisms, then evidence concerning production, recognition, recall, and language use in general can be expected (in principle) to have bearing on the investigation of rules of grammar. (1980, 200-201)
> …the biolinguistic approach focuses attention on a component of human biology that enters into the use and acquisition of language… Call it the "faculty of language"… (2005, 2)

It is clear that the FL is embedded within the performance mechanisms and operates in the use and acquisition of language. It is involved in the acquisition of language through the fixation of the parameters and also comes into play whenever a speaker/hearer processes linguistic representations in production and comprehension. The emphasis on this distinction is a methodological one. According to Chomsky, a naturalistic study "is a particular

human enterprise that seeks a special kind of understanding, attainable for humans in some few domains when problems can be simplified enough" (1995b, 10). Naturalistic linguistics accounts for a very restrictive range of phenomena whose formulation is fairly precise and whose outcomes can be empirically evaluated. Thus, the object of study is the cognitive system: its initial state (UG) and its steady state (I-language). This leaves aside the articulatory-perceptual and conceptual-intentional performance systems because they are highly complex systems in which numerous factors intervene. These factors cannot be isolated and simplified in order to study them and are, in this sense, ill-suited for a naturalistic inquiry (Chomsky 1991, 40).

In this way, the I-language is studied as a mechanism that is relatively independent from the articulatory-perceptual and conceptual-intentional systems in such a way that its explanation is not reducible to those systems. In other words, although the relative autonomy of the I-language is highly restricted by the fact that it must satisfy the legibility conditions imposed by the performance systems, there could nevertheless exist well-formed constructions that would not be usable by those systems. However, this relative autonomy of the I-language mechanism does not run counter to the fact that the I-language is a mechanism that participates in the processing of linguistic representations in language use. The distinction between competence and performance is therefore a methodological one and does not project a substantial distinction between knowledge and mechanism or between an abstract computational system and a causal computational mechanism.

As Matthews (2008) reminds us, this distinction does not allude to two parts of the mind/brain; instead, there is only one FL that is "the cognitive faculty responsible for language production and understanding" (48). The function in intension that goes from lexical items to pairs of sound-meanings is simply abstracted from concrete specification related with performance and, I would add, from structural properties related with its physical base of realization. There is a sense in which most of the models in cognitive science are formulated in these terms. The description of the mechanism of the I-language is a description of an idealized mechanism like every idealization in science. The fact that the function in intension abstracts from concrete specifications related to performance does not imply that the

grammars do not specify (in an intensional way) the pairs that the speaker/hearers compute in language processing. Although this function (pairs of sound-meanings) is not one that is computed in the processing (mappings from sound to meaning or vice versa), the first specifies the second, albeit under idealization and approximation. This specifies the domain and range of the function that is computed in the course of language production and comprehension so that the processing must respect those pairings. In this sense, "grammars *are* empirical hypotheses about psychological processes, albeit, hypotheses articulated at a level of significant abstraction" (Matthews 2006b, 465). Chomsky himself (1980) considers that grammars are specified in Marr's computational level (Marr 1982). In this sense, theories of competence or I-language would be high-level theories about psychological mechanisms (Matthews 2006a).

Of course, on the one hand, one can think that none of what I have said contradicts anything Collins says. Collins's main point is that the FL cannot be identified with mechanisms either of production or comprehension because it is involved in both; it does exactly the same work whether we are talking about production or comprehension, so it must be something more abstract. On the other hand, one can believe that Chomsky's claims hardly establish that the FL is a component of the causal mechanisms that underlie linguistic performance.

Regarding the first point, I rather like Collins's presentation of the distinction in which the FL is seen as an ensemble of systems that enter into the production and understanding of language:

> As so far characterized, competence marks out one system of an ensemble that in unison accounts for performance. In effect, then, there are no performance systems in distinction to a system of competence; there is rather, a performance ensemble including a language faculty whose properties are not exhausted by or solely dedicated to perception and production[…]. (Collins 2007, 885)

As I see it, this ensemble is the FL in a broad sense (Hauser, Chomsky, and Fitch 2002). We can consider that this broad FL includes the FL in a narrow sense as one of those mechanisms; that is, it includes a sub-mechanism that

corresponds to the abstract description of the I-language presented in section 2.1. I agree that the I-language mechanism whose function is the mapping from lexical items to <PF, LF> pairs cannot be completely identified with the production and comprehension mechanisms. However, that does not exclude the possibility that those mechanisms make use of this common resource. This conjecture does not assume that those mechanisms are one and the same. Of course, there are important differences between the mechanisms of production and comprehension, but there also have to be connections between them. One way of understanding those connections, without positing necessary dependences between those systems, is to postulate that they exploit common recourses such as a lexicon and the recursive operation merge that constitutes the I-language mechanism. This speculative picture makes room for the fact that processes both of production and of understanding should reflect the relevant properties described in the abstract specification of the linguistic theory, though not in the same way. It can also accommodate the fact that the I-language mechanism has properties that are not explained by the production and comprehension mechanisms.

Now, one could think that this picture can be feasible only if one can make sense of a processing construal of the mapping from lexical items to sound-meaning <PF, LF> pairs; even a very idealized and abstract one. This is also connected with the second point. Maybe what Collins is saying is that we cannot do so, and hence, we cannot consider the I-language as a causal mechanism. This processing construal would require conceiving of the FL as having the property of being temporal, but, according to Collins, the FL is not temporal since a derivation does not happen in time. I concede that if one tries to show that the FL is in fact a mechanism, instead of relying on programmatic assertions to the effect that it is possible to conceive it as such, one should look at the minimalist theory itself and ask how such a possibility might develop; for example, one should address the derivational vs. representational debate (see, e.g., Lasnik 2001, Uriagereka 2002, Brody 2003) and the delicate issues regarding minimalism and parsing (see, e.g., Stabler 2011, 2013). Nonetheless, as I have mentioned, I am not trying to defend the strong thesis that the FL is a mechanism, but instead a much weaker thesis, i.e., that Collins's reasons for ruling out this possibility are not compelling. In this sense, I would like to mention in the next section

some reasons for believing that the FL can be considered a mechanism that forms part of the causal architecture of the mind/brain. I will also address a way of understanding Collins's sense of 'abstract' as atemporal.

2.3. The Faculty of Language as an Abstract Causal Computational Mechanism

The minimalist program, in its recent biolinguistic approach, illustrates how the FL might work within the abstract framework seen in section 2.1. In that sense, it is a proposal about how the mechanism of the FL is composed, what its parts do, and how the organized components give rise to the function of pairing sound and meaning via the recursive 'merge' operation which combines morphemes into hierarchical constituent structures. This theory proposes a unique generative mechanism for building linguistic representations: the recursive-syntactic computational mechanism. Thus, syntax is the only generative mechanism responsible for the recursive hierarchical structure of language, whereas the phonological and semantic components are merely interpretative. This means that the theory proposes an empirical hypothesis about the workings of the FL. If it is discovered that linguistic representations can be built without making use of this generative-syntactic mechanism, the theory will be proven wrong.

In very general terms, following Marantz (2005), the mechanism operates on the lexicon via a general merger operation. The merge takes two (atomic or complex) elements, A and B, and creates a constituent, C. This is a case of external merge; if A is part of B, it is a case of internal merge, allowing for movement or displacement, which comes freely (Chomsky 2005). The syntactic constituents produced by merge are transferred at some point in the computational derivations to the interface systems. It is assumed that once a phase (i.e., the syntactic constituent or object) is transferred, subsequent operations do not refer back to it, and this results in a saving of memory. The semantic and phonological components of the FL interpret this syntactic constituent; they are not structures built on independent rules. The process of interpretation is cyclical (phase by phase) with stronger locality restrictions. This allows for, *inter alia*, mismatches between semantic and phonological structures and contextual polysemy related to semantic interpretation.

One can delve deep into the details of the model, but this description suffices for the point I wish to make, which is that it makes sense to conceive of the FL as a causal mechanism subject to empirical testing. The mechanism of the FL, as Chomsky always emphasizes, is an 'organ' on par with other biological organs such as the systems of mammalian vision, insect navigation, and many others. Although "relating mental computation to analysis at the cellular level is commonly a distant goal" (Chomsky 2005, 2), the model is intended to be well integrated into cognitive neuroscience (Marantz 2005). In this sense, its purpose, like that of most cognitive computational models, is to capture aspects (albeit abstract ones) of the causal cognitive architecture.[4] Thus, it postulates a mechanism responsible for the pairings of sound and meaning.

Mechanisms, according to Machamer, Darden, and Craver (2000), "are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions" (3). One can think of the FL's mechanism as being composed of components (i.e. the computational mechanism of I-language and the two interface mechanisms) with their activities (or functions) that are organized in such a way as to produce the phenomenon to be explained. Each of those components is, in turn, composed of entities, activities, and their organization. Thus, the computational mechanism of the I-language is composed of entities (lexical items) that engage in activities (they merge); the organization of those entities and activities produce the phenomenon of the pairings of sound and meaning. I could further argue that the FL satisfies most of the mechanist criteria for being a mechanism, but this is not my objective. What I have tried to point out so far is that there is room in the abstract specification of the FL for conceiving of the possibility that it could be a mechanism.

This being said, I think that there is a reasonable objection to conceiving of the FL as a mechanism which merits attention. This objection states that

---

[4] Although these aspects may only have to do with the inputs and outputs, and not with the processes themselves, the models are still intended to depict the causal cognitive architecture. As we will see in section 3, cognitive models are usually componential functional models whose elements and operations map the causal processes of the mechanism modeled, albeit in a highly indirect way.

since the entities in a mechanism have structural properties such as location and orientation, and the activities must have a temporal order, rate, and duration, cognitive systems are therefore not mechanisms. One could argue in favor of the position that activities in cognitive mechanisms in particular have those properties, but as is clear from the descriptions presented in sections 2.2 and 2.3, there is no clue in the FL model about the structural properties of the mechanism beyond the assumption that it must be physically realized. The entities—i.e., syntactic objects—are individuated only in terms of their functional/causal roles, leaving aside the details of the functions' implementation. This is, I think, a way of understanding what Collins maintains when he affirms that the FL is abstract in the sense that is not temporal, i.e., that a derivation does not happen in time. The temporal dimension is one of the structural properties of a mechanism. But, the theory of the FL does not describe a process in time and in general does not describe any other structural property, so it is natural to conclude that the FL is not a mechanism.

What I would like to address in the next section is the sense in which the entity responsible for the linguistic phenomena to be explained can be considered a mechanism despite the fact that the model does not describe its structural properties (3.1); and also why, in contrast to the mechanistic approach, despite the fact that there are no good reasons to rule out the possibility that it could be a mechanism, its explanation is not mechanistic (3.2).

3. Functional Models and the Mechanism of the Faculty of Language

One can read the inference mentioned in sections 1 and 2.2, which goes from the fact that the theory of the FL is about the abstract aspects of the mind/brain to the conclusion that it does not describe a causal mechanism, as stating that since the theory only describes functional properties of the entity responsible for the cognitive phenomena to be explained—leaving aside structural properties, such as temporal organization—, this entity cannot be a causal mechanism. This inference, as has also been advanced in sections 1 and 2.2, seems to be a case of the inference held by some advocates of functional explanation in cognitive science against mechanistic explanation. The defenders of functional explanation consider that

psychological explanation does not capture mechanisms (Fodor 1968, Cummins 1983, Weiskopf 2011). Their argument is based on the fact that since psychological explanation, which makes use of cognitive models, is not a mechanistic but rather a functional explanation, it does not account for mechanisms. I would like to discuss the sense in which the FL model can be seen as a functional model of a cognitive mechanism (3.1) and the reasons why, despite the fact that that there are no good reasons to rule out the possibility that it could be a mechanism, its functional explanation is not a kind of mechanistic explanation (3.2).

3.1 Functional Models of the Cognitive Mechanism of the Faculty of Language

According to the defenders of functional explanation, cognitive models intend to account for a cognitive capacity in terms of a system's functional properties and their organization. In order to do so, these models postulate a representational system, a relevant set of processes run by that system, and available resources that interact with the system's operations. The cognitive capacity in question is explained accounting for how the system is organized in terms of a sequence of sub-processes that deliver outputs on the basis of its inputs. Cognitive models are a kind of functional analysis or explanation (Weiskopf 2011). The explanandum of the functional explanation, like that of the mechanistic explanation, is the fact that a system S has the capacity for F (Cummins 1983), and this capacity is explained in terms of its functional properties and their organization. Following Piccinini & Craver (2011), three kinds of functional explanation can be distinguished in the literature about psychological explanation, depending on the kind of functional properties they make use of. *Task analysis* consists of the decomposition of a capacity into sub-capacities and their organization; in this case, the capacity for F is not explained in terms of the capacity for G of the parts of S, but instead in terms of the activity of S itself (Fodor 1968, Cummins 1983). The *functional analysis by internal states* describes the organization of those states and their interactions (Fodor 1968). Finally, *boxology* appeals to the decomposition of processes and the organization of functional individuated components (Fodor 1968).

Now, the FL model specifies the capacity of the cognitive system of the
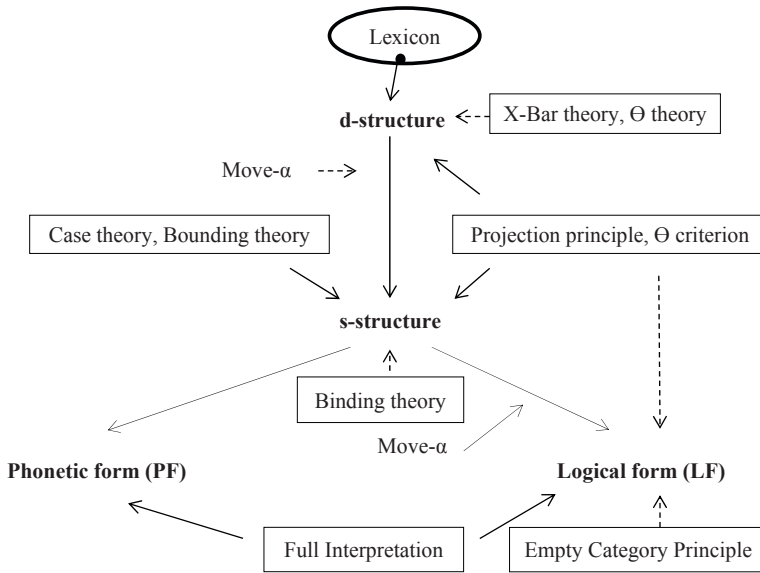
Fig. 1. Government-binding theory. Extracted from Mukherji (2010), p. 37.

FL for pairing sound and meaning in terms of its functional properties. A system's functional properties consist of the "activities or manifestations of their causal powers, dispositions, or capacities" (Piccinini & Craver 2011, 291). What remains unclear is the kind of functional model to which the FL model would belong. If we think of one of the generative models previous to the minimalist program, such as government and binding theory (Chomsky 1981), which is articulated around a sub-component related to the system of grammatical rules and a sub-component of principles, it seems to be a boxological model. In the first sub-component, a lexicon, the d-structure and s-structure, and the components of PF and LF can be distinguished. The second sub-component includes the principles of bounding theory, government theory, θ-theory, binding theory, case theory, and control theory (see figure 1). The lexicon, d-structure, s-structure, PF, and LF would be the components functionally individuated in virtue of the outputs that they produce given certain inputs, and the arrows, as in any boxological model, would reflect the functional/causal relations.
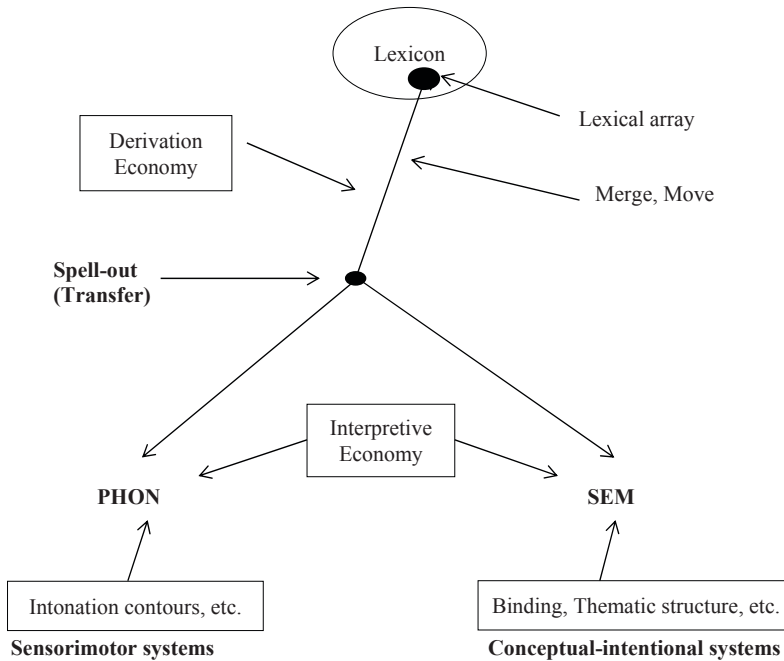
Fig. 2. Minimalist program. Extracted from Mukherji (2010), p. 164.

However, it is not clear that the minimalist model (Chomsky 1995a) described in sections 2.2 and 2.3, which is motivated by principles of minimal and optimal design, can be shaped in terms of a boxological model. Nevertheless, in favor of the *boxological* interpretation, it can be noted that like boxological models, there are components (lexical items), processes that operate upon this linguistic representations (merge), and available resources that interact with the operations (*inter alia*, principles of economy). In this sense, the capacity in question is accounted for in terms of the organization of functionally individuated components; in other words, in terms of how the system is organized in a sequence of sub-processes that produce outputs from its inputs (see figure 2).

However, it is also possible to interpret the minimalist model in terms of *functional analysis by internal states*. Given the approach of the I-language as a sequence of internal states of the speaker/hearer's mind/brain, these

would be individuated by their functional properties, i.e. by their relations with their inputs, outputs, and other internal states. In this sense, the initial state of the FL (the UG), the intermediate states of the linguistic processing, and the steady state attained (the I-language) would be explained by the internal states of the components of the computational mechanism. If we are considering the functional specification of the FL in terms of a program that is represented or run on a computer (Chomsky 1983, 201), it is natural to think of its internal states as states of its components (e.g., of its processor or memory). Thus, the explanative burden would fall on the functionally individuated components, in the same way as in boxology.

If we take into account Chomsky's assertions regarding the theory of the FL, however, it seems that he is thinking of a functional explanation in terms of *task analysis*. Accordingly, he affirms:

> We may impute existence to the postulated structures at the initial, intermediate, and steady states in just the same sense as we impute existence to a program that we believe to be somehow represented in a computer or that we postulate to account for the mental representation of a three-dimensional object in the visual field. (1980, 188)

In this sense, the quote would fit nicely into Cummins's description of psychological explanations as property theories (Cummins 1983). The aim of a property theory is to answer the question: ¿In what consist for a system S to have the property P (or, more precisely, to instantiate the property P)? In order to provide an answer, this theory appeals to properties of the components of S and how they are organized. The main idea is to explain a complex cognitive capacity of information processing (complex dispositional properties) by decomposing the system into sub-capacities that are simpler and less sophisticated and problematic than the capacity we aim to explain; such that the manifestation of the analyzing capacities (specifiable in a program) equate to the manifestation of the analyzed capacity. Thus, the first step in functional analysis consists of the analysis of the cognitive capacity P, and the second, of accounting for the instantiation in S of the analysis of P. If S possesses the components organized in the way specified in the analysis, then S possesses P. Functional explanation concludes when S runs

the program specified by the analysis (S must have the structure required for carrying out the program's instructions). In this way, all the explanatory power of the functional explanation would lie, ultimately, in the specification within the program of the analyzing capacities.

We can consider that when a program that runs in a computer is described, this is not a description of any component of the computer. In this sense, the description does not give any information about the underlying mechanism. However, Chomsky's quote above continues as follows:

> Evidence bearing on empirical hypotheses such as these might derive from many and varied sources. Ultimately, we hope to find evidence concerning the physical mechanisms that realize the program[…]. (1980, 188)

In this sense, when we talk about a computer that produces certain behavior in virtue of running a program, this means that its states are parts of some of the computer's components, because, if they were not physically instantiated, the program could not be run, and therefore the behavior in question could not be generated. Chomsky's affirmation alludes to the idea that the theory of the FL is an abstract specification of a *physical system*: the brain.[5] Thus, as Piccinini asserts: "an abstract description of a physical system is not a description of an abstract object but rather a description of a concrete system that omits certain details" (2010, 11).

Whatever the interpretation of the kind of functional explanation the FL consists of, it seems inevitable that it be committed, at least, to functionally individuated components. As we have seen, the components or parts of a mechanism have both structural properties (e.g., localization, form, orientation, and temporal organization of the sub-components) and functional properties (activities or manifestations of their causal powers, dispositions,

---

[5] "When I use such terms as 'mind,' 'mental representation,' 'mental computation,' and the like, I am keeping to *the level of abstract characterization of the properties of certain physical mechanisms*, as yet almost completely unknown" (Chomsky 1980, 13, emphasis added). In more explicit terms, Chomsky affirms: "[T]here is a special *component of the human brain* (call it 'the language faculty') that is specifically dedicated to language" (1996, 13, emphasis added).

or capacities). In this sense, the parts or components of a mechanism can be identified by their structural properties, such as spatio-temporal organization, size, etc., as well as by their functional properties in terms of their functional/causal roles.

According to Weiskopf (2011), even though the components in cognitive models are functionally individuated so that the sub-capacities are organized in a componential way, these components are not real parts of the system. Cognitive models "often posit elements that have no mechanistic 'echo': they do not map onto parts of the realizing system in any obvious or straightforward way" (Weiskopf 2011, 332). Thus, there can be neural mechanisms underlying cognitive models, but cognitive models do not capture those mechanisms. What they capture is a functional level of abstraction; they are not models of real-world mechanisms, and therefore "psychological models do not describe mechanisms" (Weiskopf 2011, 336). The components, activities, and organization that appear in functional models (unlike those that appear in mechanistic models), according to Weiskopf (2011), do not have a direct relation of correspondence to the components, activities, and organization of the system in which they are realized.

The relationship is highly indirect in the sense that "one cannot in any simple or straightforward way read off the presence of the higher level state from the lower level state." (Weiskopf 2011, 328). The nature of the mapping relationship between functional and structural properties is usually opaque, because either there is no well-defined physical system to decompose into parts, or we have no indication of what parts or operations must be used in the decompositions. Nonetheless, there are techniques that allow us to offer indirect accounts.[6] Functional models that use these techniques, according to Weiskopf, can be perfectly adequate or legitimate even if on the neurophysiological level there are no entities with the characteristics

---

[6] Weiskopf (2011) mentions reification (for example, postulating entities with the characteristics of a stable and durable object when in fact there is no such thing), functional abstraction (decomposing a system in virtue of what its sub-systems do, and not in virtue of their correspondence with the organization of the system in which they are realized), and fictionalization (including in the model components that we know do not correspond to any component in the modeled system, but fulfill an essential role for the model to be operative).

postulated by the model. What matters in order to qualify as an adequate model is to satisfy the usual criteria of adequacy for explanations.[7]

Thus, cognitive models differ from mechanistic models in the way in which they correspond to the modeled system. Whereas every element that appears in a full-blown mechanistic model corresponds to real parts of the mechanism, in functional models, according to Weiskopf (2001), "not everything that counts as a component from the point of view of the *model* will look like a component in the modeled *system* itself—at least not if our notion of a component is based on a distinct, relatively localized physical entity like a cortical column, DNA strand, ribosome, or ion channel" (332). An echo of this argument can be found in Collins (2006) when he argues that there are clearly states and processes that realize the I-language and that the theory of the FL is, in a certain degree of abstraction, about those processes; nevertheless, "we don't imagine a clean map from posited elements or operations of the faculty and physical processes–there is no tape in the head with symbols on it" (498).

Now, Weiskopf's argument, which we have just seen, can be stated as follows: if the components and activities of the cognitive model correspond to components and activities of the modeled system in the indirect ways mentioned above, it follows that cognitive models do not capture the mechanistic structure of 'real' parts, and hence do not describe mechanisms. This argument supposes a strong link between the notion of 'mechanism' and 'distinct physical localization' of real parts. Examples of real components include those already mentioned: a cortical column, DNA strand, ribosome, etc. Nonetheless, as has also been mentioned, definitions of mechanisms can be offered in functional or structural terms. In fact, Glennan affirms that a mechanism is "a complex system that produces that behavior by the interaction of a number of parts, where the interactions between parts can be characterized by direct, invariant, change-relating generalizations" (2005, 445).

Bechtel and Richardson (1993) offer a distinction between 'localization,'

---

[7] These include evidential support, representational accuracy, genuine explanation (not merely descriptive or predictive), and plausibility and consistency with the rest of our scientific knowledge.

i.e. the identification of the mechanism's parts based on their spatial localization, and 'decomposition,' i.e. the identification of the mechanism's parts based on their functional relations. With this distinction in mind, we can see that Glennan's definition is expressed in terms of functional relations among parts. Thus, according to Glennan, "it is the functional structure revealed by decomposition that is constitutive of the mechanism" (2005, 447).[8] On the other hand, there can also exist definitions of mechanisms that emphasize their structural properties. Therefore, the components or parts are simply the entities of the physical mechanism. For example, if we are trying to account for the mechanism of action potential, the components of the mechanism will be sodium and potassium channels, ions, and protein chains (Craver 2006). Whereas the structural components are specified using anatomical techniques that allow for the characterization of different brain regions in virtue of the kind of neurons and their relation, the functional components are specified in terms of their effects under certain circumstances.

Furthermore, according to Piccinini & Craver (2011), the idea of real parts in a structural sense does not imply that "the components involved are neatly spatially localizable, have only one function, are stable and unchanging, or lack complex or dynamic feedback relations with other components. Indeed, a structural component might be so distributed and diffuse as to defy tidy structural description[…]" (291). In this same line, most mechanists do not defend isomorphic correspondence, but rather maintain that the models that represent mechanisms hold similarity relations in various degrees and respects (Glennan 2005). Nonetheless, Weiskopf (2011) considers that weakening the constraint on spatial organization so that it could be distributed in parts that are not strictly localizable renders the notion of 'mechanism' empty, since it makes the task of identifying mechanisms and their particular characteristics impossible. I believe that the question about strict spatial organization as a fundamental criterion for the identification of mechanisms cannot be settled without empirical evidence. In fact, recent empirical evidence suggests that many cognitive functions are localized in vast areas. However, this is not an obstacle for current scientific research.

---

[8] Another definition of "mechanism" in completely functional terms can be found in Piccinini (2008), among others.

Therefore, it does not seem that this evidence deprives the notion of mechanism of content, provided that there are a set of criteria that can serve as a guide to delimit mechanisms.

In summary, in contrast to Weiskopf's claim, there can be psychological explanations of cognitive mechanisms that are completely functional, in that they only describe the functional properties of the mechanism, abstracting from structural properties. Weiskopf's insistence that psychological explanations do not capture mechanisms goes hand in hand with his idea that 'mechanism,' in the context of this debate, is synonymous with 'neural mechanism'. Thus, affirmations like the following can be understood: "there may be an underlying neural system, but this mechanistic structure is not what cognitive models capture" (2001, 333). This kind of statement is quite similar to Collins's assertions regarding the FL. Collins (2006) affirms that "linguistics is not in the business of characterizing a physical process" (497), and that "the grammar is not a description of any physical process, but the output of the grammar must be structured in such a way that a physical process could generate the structures[…]" (498). Thus, according to Collins, even though there are states and processes of the brain that realize linguistic competence, and the theory of the FL is, in some way, about those processes since it is an abstract characterization, it does not capture a mechanism. But to insist on the point already stressed, 'mechanism' is not synonymous with 'structural properties.' A neurobiological mechanism, such as the mechanism of the action potential or any other physical mechanism, also has functional properties, and cognitive mechanisms have structural properties, as well (though most of them are unknown).

Thus far, there do not seem to be good reasons to consider, as the defenders of functional explanation maintain, that functional abstractions do not capture real components or properties of a (cognitive, neural, etc.) mechanism. Consequently, there would not be conclusive reasons to hold that the theory of the FL, understood as a functional model, does not describe a mechanism. As I mentioned in section 1, my intention is not to show that this architectural conception is in fact the perspective that Chomsky endorses, but rather that this version fits nicely with biolinguistic developments. In this sense, Chomsky always remarks that he considers the FL to be the name of a mind/brain mechanism that linguists study indirectly,

describing its functional properties in the hopes that we can better understand its structural properties as neurobiological studies advance in their investigations (Chomsky 1980; 188, 211).

3.2 Functional and Mechanistic Models of Cognitive Mechanisms

Now, the fact that the entity responsible for the production of the cognitive phenomena to be explained may be conceived as a mechanism does not imply, contrary to what defenders of the mechanistic explanation argue, that the explanation must be mechanistic. Piccinini & Craver (2011) affirm that "Mechanistic explanation is the explanation of the capacities (functions, behaviors, activities) of a system as a whole in terms of some of its components, their properties and capacities (including their functions, behaviors, or activities), and the way they are organized together" (291). This kind of explanation involves accounting for a set of phenomena by identifying the mechanism that produces them. The description of the mechanism responsible for the phenomena should figure into the explanans. In this way, the explanans is a model of the components, activities, and organizational features of the mechanism underlying the phenomenon we wish to account for.

Piccinini & Craver (2011) believe that psychological explanation is a kind of mechanistic explanation. Since functional properties are an aspect of mechanistic explanation, functional explanation cannot be autonomous, but rather constitutes an elliptic form of mechanistic explanation (what it is called a 'mechanism sketch') which lacks some structural aspects that, when filled out, turn the explanation into a complete or full-blown mechanistic explanation. According to this conception, since "the *target* of the description in each case is a mechanism" (Piccinini & Craver 2011, 290), whether we place emphasis on the description of its structural properties or on that of its functional properties, the explanation is mechanistic.

However, following Weiskopf (2011), if what determines that functional and mechanistic models are different kinds of models is the direct or indirect relation of correspondence among the postulated components, activities, and organization of the model with those of the modeled system, respectively, then one could maintain that although the responsible entity is a mechanism, provided there are indirect relations between the cognitive model and the mechanism, the explanation cannot be considered mecha-

nistic. Similarly, one could believe that the fact that the FL is a mechanism does not imply that the explanation offered by the theory of the FL is a mechanistic one; that is, a functional explanation which lacks structural details (a mechanism sketch) that becomes a full-blown mechanistic explanation once those details are filled out. I think that Collins would agree on this point. In terms of the debate between functional and mechanistic explanation, what I am saying is that the sort of computational description provided by linguistic theory of the FL is a kind of functional description that would *not* support a causal mechanistic interpretation, given the indirect relations between the cognitive model and the FL mechanism.

Very often, components, activities, and organizations that appear in cognitive models do not have a direct correspondence to the components, activities, and organization of the physical system in which they are realized. They relate in an indirect way, in the already quoted sense that "one cannot in any simple or straightforward way read off the presence of the higher level state from the lower level state" (Weiskopf 2011, 328). If there are neatly localized parts, then it is possible, as in many cases in neurobiology, to dissect the mechanism into parts through observational and experimental methods. When there are no neatly localized parts, such as in high-level cognitive mechanisms, the mechanism can be identified, and its parts can be studied through indirect inference methods that test the processes' temporal organization or vary the (precipitating, inhibiting, background, etc.) conditions of the phenomena to be explained (Glennan 2005).

This is what Chomsky (1980, 190-191) seems to suggest when he addresses the question of the indirect empirical evidence that can be obtained in relation to the theory of the FL. He does so using an analogy about the problem of determining the nature of the thermonuclear reactions in the interior of the sun. Even though the available techniques only allow us to study the light emitted by the exterior layers of the sun, scientists can make conjectures about the fusion of light elements that form heavier elements which convert mass into energy, and in this way account for hidden thermonuclear reactions. The evidence offered by the light emissions from the periphery of the sun constitutes indirect evidence. There is no way to enter the interior of the sun to obtain direct evidence. Similarly, the inquiry into the FL makes use of indirect evidence, such as reaction time, recogni-

tion, informants' judgments, etc. The difference in the analogy between the postulation of physical entities and processes, in the case of the sun, and the specification abstracted from structural properties of the brain, in the case of language, is not a principled one, however. For ethical reasons, we simply cannot explore the mechanisms by direct experimentation, but this does not generate any particular problem for the investigation of cognitive capacities (in any case, this would be analogous to the case of astronomers who can only study the light emissions from the periphery of the sun).

In this way, the kind of explanation does not depend on the kind of entity responsible for the phenomenon to be explained (whether or not it is a mechanism), but rather on the kind of epistemic access to that entity that is available to us. If the correspondence or mapping relationship among the postulated entities, activities, and organization in a model is directly related to the parts, activities, and organization of the mechanism, there are good chances of achieving a mechanistic explanation. This is what is expected in general in the cases of multiple levels of neurobiological mechanisms or low-level cognitive mechanisms. If, in contrast, the correspondence is indirect, as in the case of high-level cognitive mechanisms, a functional explanation seems more appropriate to the extent that it highlights the relevant explanatory characteristics of the mechanism. This does not necessarily mean that a functional explanation is a mechanism sketch such that when the structural aspects are 'filled out,' it becomes a full-blown mechanistic explanation. Mechanistic sketches are precisely the kind of explanations which are considered capable of being filled out by structural details. However, in the case of the FL, perhaps most of these details cannot be filled out for reasons of epistemic access, simply because "one cannot in any simple or straightforward way read off the presence of the higher level state from the lower level state" (Weiskopf 2011, 328).

4. Conclusions

In this paper, I have tried to dismantle the inference that goes from the affirmation that the theory of the FL is an abstract specification of the mind/brain to the conclusion that its object of study, the FL, cannot be conceived of as a cognitive causal mechanism. I believe that this inference in the field

of the philosophy of linguistics is a case of a more general inference that is widely spread in the philosophy of cognitive science by defenders of functional explanation. In this sense, I have tried to defend the idea that it is perfectly acceptable to maintain that a certain theory or model only offers an explanation that alludes to functional properties of the entity responsible for the cognitive phenomena we wish to explain; while its object of study can be conceived of as a mechanism that possesses both structural properties and functional properties that manifest causal dispositions, allowing it to be inserted in a causal nexus. In consequence, there is no need to reject that psychological explanation captures mechanisms in order to defend some form of abstract specification or functional psychological explanation, as I think Collins and Weiskopf do. Nor is it necessary to insist that psychological explanation is mechanistic because it accounts for mechanisms, as I think Piccinini & Craver do.

Maybe the difference between the abstract and the causal architectural version of the FL is simply a difference of emphasis. Both versions agree that the specification or explanation of the FL is couched in functional or abstracts terms in the sense that this specification does not capture structural properties, such as a temporal dimension. The difference apparently lies in the fact that, from the causal architectural perspective, although the description of the FL does not capture its structural properties, it does describe a physically instantiated cognitive mechanism that has functional as well as structural properties. The abstract architectural version seems to be more cautious or skeptical about such possibility. Assuming that the FL is instantiated in a physical mechanism, I suppose that the overall picture of the highly indirect relations between the cognitive model and the cognitive mechanism will be shared by both the abstract and the causal architectural perspectives.

References

Baker, M. 2001. The Atoms of Language. The Mind's Hidden Rules of Grammar. New York: Basic Books.

Bechtel, W., & Richarson, R. 1993. Discovery Complexity: Decomposition and Localization as Scientific Research Strategies. Princeton: Princeton University

Press.

Brody, M. 2003. Towards an Elegant Syntax. London: Routledge.

Chomsky, N. 1965. Aspects of the Theory of Syntax. Cambridge, MA: MIT Press.

Chomsky, N. 1981. Lectures on Government and Binding. Dordrecht: Foris.

Chomsky, N. 1980. Rules and Representations. New York: Columbia University Press.

Chomsky, N. 1991. Linguistics and Cognitive Science: Problems and Mysteries. In A. Kasher (Ed.), The Chomskyan Turn, 26-53. Oxford: Basil Blackwell.

Chomsky, N. 1995a. The Minimalist Program. Cambridge, MA: MIT Press.

Chomsky, N. 1995b. Language and Nature. Mind, 104(413), 1-61.

Chomsky, N. 1996. Powers and Prospect: Reflections on Human Nature and the Social order. London: Pluto Press.

Chomsky, N. 2000. New Horizons in the Study of Language and Mind. Cambridge: MA, MIT Press.

Chomsky, N. 2005. Three factors in language design. Linguistic Inquiry 36(1), 1-22.

Collins, J. 2004. Faculty Disputes. Mind & Language 19(5), 503-533.

Collins, J. 2006. Between a Rock and a Hard Place: A Dialogue on the Philosophy and Methodology of Linguistics. Croatian Journal of Philosophy 6, 471-505.

Collins, J. 2007. Linguistic Competence without Knowledge of Language. Philosophy Compass 2:6, 880-895.

Collins, J. 2008. Knowledge of Language Redux. Croatian Journal of Philosophy 8(22), 3-42.

Collins, J. 2009. The Perils of Content. Croatian Journal of Philosophy 9(27), 259-289.

Craver, C. 2006. When Mechanistic Models Explain. Synthese 153(3), 355-376.

Cummins, R. 1983. The Nature of Psychological Explanation. Cambridge, MA: MIT Press.

Davies, M. 1986. Tacit Knowledge and the Structure of Thought and Language. In C. Travis (Ed.), Meaning and Interpretation, 127-158. Oxford: Oxford University Press.

Dwyer, S, & Pietroski, P. 1996. Believing in Language. Philosophy of Science 63(3), 338-373.

Evans, G. 1981. Semantic Theory and Tacit Knowledge. In S. H. Holzman & C. M. Leich (Eds.), Wittgenstein: To Follow a Rule, 118-137. London: Routledge.

Fodor, J. 1968. Psychological Explanation. New York: Random House.

Fodor, J. 1975. The Language of Thought. Nueva York: Harper & Row.

Fodor, J. 1983. The Modularity of Mind. Cambridge, MA.: MIT Press.

Fodor, J. 2001. Doing Without What's Within: Fiona Cowie's Critique of Nativism. Mind 110(437), 99-148.

Garrett, M.F. 1982. Production of Speech: Observations from Normal and

Pathological Language Use. In A. Ellis (Ed.), Normality and Pathology in Cognitive Functions, 19-76. London: Academic Press.

Glennan, S. 2002. Rethinking mechanistic explanation. Philosophy of Science 69, 342-353.

Glennan, S. 2005. Modeling mechanisms. Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences 36, 445-464.

Hauser, M. D., Chomsky, N., and Fitch, W. T. 2002. The Faculty of Language: What Is It, Who Has It, and How Did It Evolve? Science 298, 1569-1579.

Knowles, J. 2000. Knowledge of Grammar as a Propositional Attitude. Philosophical Psychology 13(3), 325-353.

Lasnik, H. 2001. Derivation and representation in modern transformational syntax. In M. Baltin and C. Collins (Eds.) Handbook of syntactic theory, 62-88. Oxford: Blackwell

Machamer, P., Darden, L. & Craver, C. 2000. Thinking about Mechanisms, Philosophy of Science 67, 1-25.

Marantz, A. 2005. Generative linguistics within the cognitive neuroscience of language. The Linguistic Review 22, 429-445.

Marr, D. 1982. Vision. San Francisco: W. H. Freeman.

Matthews, R. 2006a. Knowledge of Language and Linguistic Competence. In E. Sosa, & E. Villaneuva (Eds.), Philosophy of Language: Philosophical Issues 16, 200-220. Oxford: Blackwell.

Matthews, R. 2006b. Could Competent Speakers Really Be Ignorant of Their Language? Croatian Journal of Philosophy 6(3), 459-69.

Matthews, R. 2008. Epistemic Heresies: Reply to John Collins' 'Redux'. Croatian Journal of Philosophy 22, 45-55.

Mukherji, N. 2010. The Primacy of Grammar, Cambridge, MA.:MIT Press.

Piccinini, G. 2008. Computers. Philosophical Quarterly 89(1), 32–73.

Piccinini, G. 2010. Computation in Physical Systems. The Stanford Encyclopedia of Philosophy (Fall 2010 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2010/entries/computation-physicalsystems/>.

Piccinini, G., & Craver, C. 2011. Integrating Psychology and Neuroscience: Functional Analyses as Mechanism Sketches. Synthese 186, 283-311.

Skidelsky, L. Forthcoming. Competencia lingüística, actitudes proposicionales y representaciones. In A. Audano, C. Carrió, A. Gonzalo, G. Parera, C. Reynares (Eds.), El camino desde Syntactic Structures. Un recorrido histórico y filosófico por la lingüística chomskiana. Santa Fe: Universidad Nacional del Litoral.

Skidelsky, L. 2006. Personal-Subpersonal: The Problems of Inter-level Relations. Protosociology. Special Issue: Compositionality, Concepts and Representations II: New Problems in Cognitive Science 22, 120-139.

Stabler, E. 2010. Computational perspectives on minimalism. In C. Boeckx (Ed.)

Oxford Handbook of Linguistic Minimalism, 617-642, Oxford: Oxford University Press.

Stabler, E. 2013. Two models of minimalist, incremental syntactic analysis. Forthcoming in Topics in Cognitive Science.

Uriagereka, J. 2002. Derivations: Exploring the Dynamics of Syntax. London: Routledge.

Weiskopf, D. 2011. Models and mechanism in psychological explanation. Synthese 186, 313-338.