



Scuola Normale Superiore

Anno Accademico 2010–11

Tesi di Perfezionamento in Discipline Filosofiche

Holism You Can Buy

What Semantic Holism Is and What We Should Think of It

Relatore Prof. Massimo Mugnai

Candidato Alessandro La Grua

Preface

Come il fiume fra i monti
ha anfratti e golfi e seni
in cui nasconde liquidi sentieri,
anche la lingua cela in orizzonti
riposti ed alieni
amori, odi e pensieri.

To write this work took me relatively little. The process preceding the writing of it, though, was a tough and troubled travel, which began with my first year at the University of Pisa, when I decided that language would have been my privileged interest, and carried on during the next eight years.

My interest in semantic holism has many roots: so many different authors seemed to me to be endorsing some form or the other of holism, that it became difficult for me to see that it is indeed a controversial theory at all. Since it is my conviction that our most rooted ideas, especially when they get to lead our thoughts and reflections, are the ones we must investigate with the greatest care (hence my idiosyncrasy against arguments appealing to intuition, which may become apparent in the present work), I decided to turn myself to this topic.

I found it to be true that in philosophy a great portion of work consists in asking the right questions, rather than finding the right answers. In fact, what may look like a plausible answer today, tomorrow may be found to be utterly false, and maybe again plausible the day after. This is because philosophy does not operate in isolation, and discoveries and suggestions from other disciplines may shed a light on our investigations. Questions, on the other hand, can be intriguing or futile, but are never right or wrong, and are therefore less prone to change with time. On the other hand, good questions and good answers have this in common, in philosophy: that both can open more roads than they close, thus bringing a whole new perspective onto an ancient matter.

I do not claim to have found such good questions, nor good answers, for that matter. But I began this investigation still with a prejudice in favour of holism, and my aim was to find a battery of arguments to prove it right against its opponents. A thorough investigation of those very opponents, and of arguments in favour of holism itself, showed me how misguided I was.

Still, conversations with linguists and philosophers convinced me that I was not the only one taking semantic holism for granted. Thus I found my starting point: even if holism cannot be proved absolutely right, it may be proved more useful than its opponents from a heuristical point of view – in virtue, that is, of its consequences with respect to other disciplines such as linguistics. Although even this rather weaker hypothesis would eventually prove wrong (I found semantic atomism and molecularism to have as much heuristic value as semantic holism), to investigate it was at least within my possibilities.

Acknowledgments

This thesis would not have been possible without the support and the suggestions of many people. I will try to pay part of my debt toward them by struggling to at least remember – if not list – the name of everyone: a task which is unlikely I will fulfill.

Professor Mugnai is not only the advisor of this dissertation: he is also, thanks to his criticism to my sometimes obscure reasonings, the main responsible for the shift in direction of the thesis itself. Of course, any mistake in that path is only my fault.

Professors Bernini, Bertinetto, Higginbotham, Marconi are also essentially linked to my work. It is only because of their careful readings and criticisms that I could shape my convictions and give a written form to them. There have been contrasts and disagreements, of course, but at least they made it possible for my work to be worthy of disagreement.

Everybody in the Ph.D. program in Philosophy, here at SNS, had sooner or later to listen to my rants about semantic holism, and the principle of charity, and human rationality. I know that I was not as careful as a friend and as a listener as they were: I apologize.

A similar apology goes to the wonderful people in the graduate program of the School of Philosophy at USC, who not only treated me as one of them from the very first day I was in Los Angeles, but provided many a good idea and many a valuable criticism. Also, it was fun to be there, wasn't it?

Friends here in Pisa and abroad provided that oblique but rigorous look at the issue that is necessary not to get stuck in the conviction of being always right. For this let me thank, among the others, Valentina Bambini, Marta Ghio, Martina Johnson, Marco Santi, Emanuele Spadaro. The moral support of other friends was no less appreciated than scientific one: thus I thank Cesare Bartolini, Guido Bini, Ida Campeggiani, Gianmario Erriquez, Alberto Godioli, Fabio Guidetti, Giacomo Macrì, Filippo Mariano, Giuseppe Mazzaglia, Massimo Palermo, Alessio Panetti, Gianpaolo Pecchia, Eugenio Refini, Fabrizio Salomone, Mauro Scarabelli, Lucia Sipala, Valentina Sturli, and a plethora of people I am certainly leaving out.

Professors Acquaviva, Goldsmith, and Moro came to Pisa in various occasions to give lectures, and were so kind to give me suggestions and precious insights. Also, conversations with Dr. Carlo Marletti had great value to understand some connections and implications of my ideas.

It is with the greatest gratitude that I think of my mom and dad, who were more supportive than one can desire, and of Chiara, who was as spurring and encouraging as only future family can be.

Finally, I let my thoughts go to two great people who are no more. Professor Gargani was the most helpful teacher I could desire when we met and after. Nothing of what I think and do would be as it is, without the constant confrontation with his own opinion, which I deeply miss. Davide Soragna was a sensitive friend, whose freedom of thought taught me always to struggle for rigour more than any academic discussion with professors and colleagues did.

Contents

Preface	iii
Acknowledgments	iv
Introduction	1
1 Preliminary Notions	5
1.1 Utterances, Sentences, Propositions	5
1.2 Language	7
1.3 Reference, Meaning, Truth	13
I Competitors	17
2 Semantic Atomism	19
2.1 What Words Mean	19
2.1.1 Starting from References	20
2.1.2 Starting from Meanings	24
2.1.3 Starting from Truth	28
2.2 What People Mean	31
2.2.1 Interpretation by Comparison	32
2.2.2 Interpretation without Clue	35
2.3 Evaluating Semantic Atomism	36
2.3.1 Internal Problems	36
2.3.2 Interactions	40
2.4 Some Brief Conclusion	42
3 Semantic Molecularism	45
3.1 The Meaning and Sense of Sentences	46
3.1.1 Meaning As Truth Conditions	47
3.1.2 Meaning As Verification	50
3.1.3 Meaning As Falsifying Criteria	53
3.2 The Force of Discourses	55
3.3 Evaluating Semantic Molecularism	60
3.3.1 Internal Problems	60

3.3.2	Interactions	64
3.4	Some Brief Conclusion	65
II Semantic Holism		67
4	Holistic Interpretation	69
4.1	The Notion of Holism	70
4.1.1	About an Eccentric Definition of Holism	70
4.1.2	The Argument A	74
4.2	Translation and Interpretation	80
4.2.1	Some Issues about Behavioristical Theories of Interpretation	82
4.2.2	Again on the Analytic/Synthetic Distinction	84
4.2.3	Semantics As an Empirical Enterprise	90
4.2.4	Compositionality	94
4.2.5	Radical Interpretation	100
4.2.6	Nomologicity	103
4.2.7	Interpretation with Charity	106
4.3	Interpreting One's Own Sentences	116
4.3.1	An Atomistic Knowledge?	116
4.3.2	Social Speakers	117
4.3.3	A Rational System	118
5	Holistic Meaning	121
5.1	The Place for Linguistics	121
5.2	Formal Semantic Theories	123
5.2.1	The Two Components of Meaning	123
5.2.2	On the Relationship between the Linguistic and Conceptual Systems	126
5.2.3	The Linguistic System	128
5.2.4	The Conceptual System	135
5.3	Cognitive Semantic Theories	143
6	Evaluating Semantic Holism	149
6.1	Internal Problems	149
6.1.1	The Problems with Interpretation	149
6.1.2	The Problems with Holistic Semantics	159
6.2	Interactions	161
6.2.1	Philosophical Issues	161
6.2.2	Linguistics	170
6.3	Final Remarks	176
Bibliography		177

Introduction

What you are reading is something more and something less than a work about semantic holism. Trivially, it is more because semantic holism is not the only semantic theory I take into consideration to prove my point. It is more, also, because of the point I try to make in this work. I do not address the theme of semantic theories from the exclusive point of view of their claims concerning the way we can build a semantic theory for our languages. Rather, my aim is to investigate semantic theories from a metaholistic point of view, so to speak. Theories do not develop in isolation: their advocates operate in a rich theoretical and social environment, and are open to suggestion and motives potentially coming from any area of interest. My claim, starting from these assumptions, is that each of the semantic theories that we may devise have privileged interlocutors among other theories and disciplines: these privileged interlocutors are the theories and disciplines that more than others can profit from investigation in the direction suggested by the semantic theory in question, and that, on the other hand, more than others can furnish it with valuable suggestions themselves.

Thus, this work is more than merely about semantic holism in itself, because its principal focus is on the relationship between semantic holism and other philosophical and scientific theories. The focus, as I will sometimes say in what follows, is on the heuristic value of semantic holism, and of other kinds of semantic theories too. For this reason, this work is also somewhat less than a work about semantic holism. That is, it is not about semantic holism in itself, and it is certain that it does not cover every issue that is possible to think of about semantic holism. And this is not only because such a task would be impossible. Holistic theories about meaning, in one form or the other, are around since philosophers started to reflect about language. Without going into details, it is arguable that some kind of holistic theory of meaning can be found in what is assumed to be the first essay of philosophy of language, Plato's *Cratylus*.

Obviously there is no room, in this as in any other work, to cover every detail of such an enormous history. Then, of course, there is the fact that I did not choose to deal only with semantic holism, but also with what I dubbed its competitors – semantic atomism and semantic molecularism. Therefore,

the number of issues I should have been dealing with would at least triplicate. This reason for the incompleteness of my treatment of semantic holism, though, is more or less contingent: that is, if this were the only reason, then it would be possible to maintain that, were not the literature on the subject so immense, I would have covered it. But this is not the case: since, to repeat myself, my aim is not to deal with holism in itself, but only in its heuristic connections, I focused specifically on those aspects of holism, of atomism, and of molecularism that are relevant to highlight those connections. Therefore, there is much that is left out of this work. Some, inevitably, should have been not, for my ignorance is still great; but some of the omissions are indeed deliberate.

Some words on the structure of the present work will be adequate, now. I decided to begin with a chapter devoted to the introduction of some terminology. This is not because I have personal notations that the reader should be aware of, but because the traditional terminology of philosophy of language is itself rather ambiguously used by the scholars. Even if the terms are the same, the notions attached to them vary widely. I do not claim that my notions are to be preferred, but only that they are the most adequate to the arguments I will draw in what follows. It is important that the reader notices from the beginning that it is extremely plausible that at least some of my arguments are biased by the use I do of the terminology. The notions I think require clarification are those of **utterance**, **sentence**, and **proposition** (§ 1.1), that of **language** itself (§ 1.2), and those of **reference**, **meaning**, and **truth** (§ 1.3): as one can see, they are the very bricks of philosophy of language.

I devoted the first part of the work to the examination of the competitors of semantic holism, so that the latter would be described not only by itself, but also by the comparison with the alternative theories. The two chapters about semantic atomism and semantic molecularism are parallel. In each I start by describing what I call the properly semantic side of the theory, then I move to the interpretative side of the theory of meaning (i.e., the theory of interpretation that puts the theory of meaning into effect). Finally, I tried to evaluate each of the theories regarding both its internal assumptions and development, and its heuristic connections with other theories and disciplines.

The same schema has been adopted for the second part of the work, devoted to semantic holism. Here the three-fold structure is expanded in three chapters. It is also reversed in its first two parts, i.e., I start with the holistic theory of interpretation, and then I move to the holistic theory of meaning. Then I evaluate semantic holism: again, first by itself, and then by comparing it with its heuristic value.

It is important to state here, at the beginning of this investigation, that the heuristic evaluation of the three semantic theories attempted here does not arrive to claim one better than the others. Although my theoretical preference still goes to holism, and my sympathy goes to molecularism for reasons that

are not entirely supportable by theoretically sound arguments, I find that each of the three theories has its merits and its faults, and that, considered in the network of knowledge which is relevant, each of them is, at least for now, unreplaceable.

CHAPTER 1

Preliminary Notions

Every semantic theory deals with the meaning of sentences. We expect such a theory to be satisfactory in characterizing sentences in a language. In order to obtain this, we must preliminarily accomplish an equally acceptable characterization of language itself. Even before that, we must explain what we mean by ‘sentence.’ The latter task will be the topic of § 1.1, the former of § 1.2. We would also need to clarify the differences and relationships between such notions as ‘reference,’ ‘meaning,’ and ‘truth,’ which bear great importance for semantics. I will do this in § 1.3.

1.1 Utterances, Sentences, Propositions

The notion of sentence does not come in isolation: it is connected to those of utterance and proposition. This distinction is what I want to focus on now.

We can examine what we say under three points of view:

1. If we are interested in the physical instantiation of speech, then we are dealing with the **utterance**. Thus, this term always refers to a single physical entity. It is the result of a performance: it can be sound, gesture, graphic sign (writing).
2. Every utterance is, hence, the physical instantiation of a linguistic entity: the **sentence**. We can have, and typically we do have, more than one utterance per sentence. We may say that the sentence is the *type* of which individual utterances are *tokens*. If we deal with what has been said inasmuch as it is a linguistic product with linguistic features, then we are dealing with sentences.
3. The sentence is a linguistic entity: therefore, it has a meaning, it is the sign for something. This meaning, when we consider it as a metaphysical entity outside language, is what we call **proposition**. Since more sentences can mean the same, more than one sentence can correspond to the same proposition.

Propositions and sentences are, ontologically speaking, universals. We have said that the sentence can be seen as the type whose tokens are single utterances. Therefore, rigorous nominalists might prefer to develop a semantic theory coping with mere utterances. They would maintain that, formally speaking, only utterances really exist, and we talk of sentences only for practical reasons. Sentences may be the mere generalization of linguistic properties shared by utterances. This is Michael Devitt's position:

It is often convenient to talk of the objects posited by these theories as if they were *types* not tokens, as if they were abstract Platonic objects, but this need be nothing more than a manner of speaking: when the chips are down the objects are parts of the spatio-temporal world.¹

Up to this point, Devitt's Ockhamist position is absolutely legitimate. Devitt claims, though, that this nominalistic argument should make us deem language as the sum of linguistic outputs, i.e., of utterances.² Such a notion of language, though, makes the whole talk about language redundant,³ and if we cannot mention language, it is difficult to see how to define the linguistic features in virtue of which we are interested in dealing with utterances and not, say, with grunts. Devitt knows it: he adds indeed a couple of pages later that, in order to select relevant data, a theory has to rely on a preexistent classification, to establish what should count as linguistic data and what not.⁴

Nothing forbids that the preexistent classification comes from a robust philosophical theory, rather than folk linguistics, as Devitt would prefer. Being a nominalist with respect to utterances and sentences does not necessarily imply that we can talk about languages only by means of utterances: if we do not want to renounce talking of the infinite potentialities of language,⁵ we have to go beyond observation and the observable, which is always finite.

The notion of proposition is even more problematic than the notion of sentence. Propositions are ideal entities: their existence is explicitly mentioned only in the context of some philosophical theories (usually, semantic theories). Therefore, their existence is not accepted by everyone. Propositions, unlike sentences, are not linguistic entities; still they are expressible only by means of language. The point is that every feature we would attribute to propositions must be uncovered analysing the alleged effects of propositions on sentences. If the only access to the properties of propositions is through sentences, and if

¹Devitt (2006), pg. 26, emphasis in text.

²Cfr. *ibid.*, pg. 25: "On my view, a language is composed of the outputs of a linguistic competence, symbols that are governed by a system of linguistic structure rules."

³See § 1.2.

⁴Devitt (2006), pg. 27: "[...] guided by folk linguistics, we start with an intuitive idea of the domain of grammatical tokens to be studied."

⁵This is what below I call the Humboldtian characteristic of language. See § 1.2.

the only role of propositions in a theory is to explain properties of sentences, then propositions may seem redundant with respect to sentences:

There is indeed [...] the hint that corresponding to each meaningful expression that is an entity, its meaning. This idea, even if not wrong, has proven to be very little help: at best it hypostasizes the problem.⁶

It is easy to see the origin of the notions of sentences and propositions in their alleged convenience in theoretical reasoning, as it happens for all hypostases. To see it, though, is not to take a stand on the correctness of these notions. Such a stand can be taken only after a philosophy of language has shown, by means of a theory of meaning, whether these notions are justified, indispensable, and not ad hoc ones.

1.2 Language

Language, whatever our depiction of it, has the characteristic, famously expressed by Humboldt, of making “infinite use of finite means.”⁷ Indeed every language can express infinite different sentences by means of a finite lexicon and a finite number of syntactic rules. Every description of language, then, has to deal with this characteristic, which I will dub **Humboldtian characteristic**. I would like to show in what follows that the best way, if not the only, to cope with it is to consider languages relative to speakers, and to treat them as produced by rules of use mastered by speakers. Thus, to reconstruct, starting from lexicon, all the infinite sentences which are expressible in the language, we must assume that at least one of the rules we are considering is recursive.

The major alternative of such an approach is to consider languages as constituted by a lexicon plus a set of recursive syntactic rules plus a set of possibly recursive semantic rules. The two proposals are, as a matter of fact, one and the same. We must acknowledge that different speakers have different lexicons and they may seem to have different rules:

It could even happen that every speaker from the start had his own quite unique way of speaking. Something approaching this is in fact the case, of course. Different speakers have different stocks of proper names, different vocabularies, and attach somewhat different meanings to words.⁸

⁶Davidson (2001), pg. 126.

⁷Wilhelm von HUMBOLDT, *Über die Verschiedenheit des Menschlichen Sprachbaues*, Berlin 1836, pg. 122, cit. in Chomsky (2000), pg. 73.

⁸Cf. Davidson (2001), pgs. 276 f.

Differences in lexicons are the most easy to detect, of course. There could also be some slight difference between speakers as to what concerns syntax. Although it might be claimed that the core syntax is one and the same for every speaker, one would find ungrammatical some sentences that another would feel perfectly correct. There are ways of explaining this: for example, we can think, as generativists do, that there are parameters that are differently setted, and this different settings in turn can account for what is at stake here.

Moreover, Davidson seemed to allude in the preceding quotation to the possibility of having different semantic rules between speakers. While such a difference would surely account for the attachment of “somewhat different meanings to words,” we must stress that the determination of the element responsible for that difference may depend on the semantic theory we entertain. If we do not believe in a strong relationship between semantic and conceptual systems, for instance, we may deem the latter but not the former responsible for such a difference.

Another possibility to characterize a language is to identify it by means of sentences. We have indeed two possibilities, both of them quite unsatisfactory:

1. We may try to identify a language by means of the set of sentences which result *expressible* in it, i.e., sentences which result well-formed according to the language. We should deal, thus, with an infinite set of sentences; and we can manage to do it only by recursive means, i.e., by means of recursive rules. Mathematical recursion, in fact, is the only way we can cope with infinity.

Moreover, without a rule to determine when a sentence belongs to this set and when not, we are not able to distinguish among different languages. Such a rule must be a rule about the compositional features of the language; it cannot be derived from assumptions about what a sentence would mean or how it is used by someone who would utter it. Our experience of what languages can do is that we cannot say looking at what it means whether a sentence is well-formed in a language or not – hence, whether it belongs to the set or not. Indeed, even meaningless sentences can be well-formed in a language: take the classical example by Chomsky, “Colorless green ideas sleep furiously.” On the other hand, there are many ungrammatical sentences that are readily interpretable, such as: “Who did you say would take you where?”

We can only partially discriminate between allowed and forbidden sentences with the help of the lexicon: sentences using words not in the lexicon of the language may not be considered of that language. Or, we can have some syntactic theory that helps defining well-formedness in a language. However, both a lexicon and a syntactic theory are usually reconstructed a posteriori from linguistic behavior of speakers. What we do, then, when we rule some sentence out of a language, is claiming that

speakers of that language would never treat it as grammatical. And this means that we relativize languages to speakers.

2. Alternatively, we may treat a language as the set of all sentences actually ever *expressed* in it. This is obviously a finite set, though large, because in a finite time only a finite number of sentences is ever expressed in a language. Anyway, this kind of identification takes in no account (actually, it denies) the Humboldtian characteristic. Should we renounce it, we could not attribute a single language to a speaker. In fact, given a language – i.e., in the hypothesis under examination, a finite set of sentences – every sentence pronounced by a speaker not in the initial set would constitute a new language. Every speaker, thus, would speak a number of languages equal to the number of sentences produced. Hence, the notion of language would be redundant with respect to that of sentence, or even to that of utterance, if we think that no such things as sentences exist.⁹

Another possibility would be to trust common sense when it identifies languages such as English, Italian, Japanese, and so on. Common sense, however, is usually vague in distinguishing among languages and between what should be considered a language and what a dialect or a local variety. With more refined criteria than common sense, such as sociologic or historic criteria, we either do not defy vagueness, or we replace it with arbitrariness.

It may anyway result useful for an intelligible discourse to draw linguistic boundaries and to have a notion of speech community to start with. Such a notion, although formally incorrect, may constitute a valid heuristical start for scientific or philosophical enterprises.

For instance, nearly all linguists and some philosophers of language (e.g., Davidson and his followers) maintain that languages are always relative to speakers: i.e., that, strictly speaking, only **idiolects** (individuals' languages) exist. Idiolects are identified by means of rules that are specific to every speaker. Both disciplines, though, employ the notions of common language and speech community because of the heuristical power of such notions.

Obviously, the employment of such a contradictory notion in a strict theoretic elaboration must be acknowledged. This is what Chomsky and Davidson did, respectively for linguistics and for philosophy of language:

Of course, it is understood that speech communities in the Bloomfieldian sense – that is, collections of individuals with the same speech behavior – do not exist in the real world. Each individual has acquired a language in the course of complex social interactions with people who vary in the ways they speak and interpret

⁹See § 1.1.

what they hear and in the internal representations that underlie their use of language.¹⁰

Given a community of speakers with apparently the same linguistic repertoire, [...] the theorist will strive for a single theory of interpretation: this will greatly narrow his practical choice of preliminary theories for each individual speaker.¹¹

Davidson and Chomsky, then, know very well that we cannot talk about the same linguistic repertoire of different speakers. Their theories, although somehow different, enable them to distinguish among more or less similar repertoires, and to employ the idealized notion of speech community. This is all they need to pursue their discourses.

Some scholars have claimed, though, that linguists' usage of the notion of speech community is, indeed, something more than sheer pragmatism:

[...] the dismissal of concern with "a common language" like English seems overblown, at least. The members of any group that share a meaning of one linguistic expression tend to share meanings of a vast number of others and it is convenient, on the basis of this, to follow the custom of classifying sets of these expressions with shared meanings as English, Spanish, and so on. The classification is bound to be a bit vague but no more so than many scientifically appropriate ones. And such classifications *seem* to be useful in linguistics, for linguistics books and articles are replete with them. Are these classifications mere manners of speaking that can be paraphrased away when the serious linguistic work is done? I think not: they are necessary for the linguist *to identify what she is talking about, to identify the subject matter*. For, the subject matter is the shared meanings and syntactic properties of linguistic expressions *in a certain group of people*. To identify this subject matter she has to identify the group.¹²

One should note that Devitt does not claim that the notion of idiolect is dispensable: in fact, to say that linguistic research must deal with shared meanings among idiolects¹³ is the same as to say that idiolects have a primitive position, in an ontological sense, in an explanatory sense, and in a temporal sense. On the other hand, I do not see why we cannot say that the identification of groups is a theoretical construction: it is undoubtedly useful, but

¹⁰Chomsky (1986), pg. 16.

¹¹Davidson (2001), pg. 153.

¹²Devitt (2006), pgs. 183 f., emphasis in text.

¹³Cf. *ibid.* (emphasis in text): "[...] my *sixth major conclusion*: the primary concern in linguistics should not be with idiolects but with linguistic expressions that share meanings in idiolects."

it is not independent of theory. To say that the subject matter of linguistics is meanings and expressions shared among members of a group even before having investigated how to identify such groups is to put the cart before the horse.

I agree with Devitt when he says that linguistics shows efficient and coherent employment of the classification in linguistic groups, and that often linguists cannot avoid mentioning such groups. On the other hand, there are plenty of classifications that proved useful to scientific goals and that are, nonetheless, irredeemably vague and therefore, perhaps, even false, if taken at face value as ontological classification – e.g., the classification in species as it is used in zoology.¹⁴ Still, in Devitt’s phrase, they are “scientifically appropriate.” Ontological existence and priority of entities cannot be deduced from the existence of some science dealing with those entities. Theoretical constructions and scientific classifications are so debated, that every ontological deduction based on them should be discouraged.

What about artificial languages? Languages of this kind are possibly without speakers. We can conceive them, though, as languages that a ‘speaker’ could ‘speak,’ i.e., that some kind of rational being might use for its purposes. Of course, in certain cases, this would entail some odd kind of ‘speaker,’ but not so odd we might not want to talk of ‘speaking’ any more, albeit in a rather metaphorical way or, as I do, with scare quotes. In other words, I think that we decide which artificial languages count as languages by comparing them with natural ones.

It is certainly possible to conceive of an artificial language constituted by such rules and lexicon, that we are able to construct only a finite number of sentences. In this case, however, we must decide if it is still appropriate to talk of language. If we take the Humboldtian characteristic to be part of the definition of language, then a language with a finite number of possible sentences is internally contradictory. This seems to be the opinion of those who maintain recursion to be the core of the faculty of language.¹⁵

On the other hand, if we deem the Humboldtian characteristic and recursion merely as sufficient but not necessary conditions to talk of language, we must acknowledge that even a symbolic system with only one sign may count as one language, then. The possibility of this kind of **punctate** languages, as Fodor and Lepore call them¹⁶, is one of the objections that they consider to

¹⁴I am not claiming that there cannot be an ontological theory to metaphysically validate the zoological taxonomy: what I am saying is that there is no guarantee that such a theory is feasible, that theories supporting the opposite thesis have indeed been sustained, and, more importantly, that there is no desire from the point of view of biologists of furnishing a taxonomy with any ontological value, or with any more value than that which biology itself would assign to it.

¹⁵Cf. Hauser et al. (2002).

¹⁶Fodor and Lepore (1992), pg. 32.

semantic holism.¹⁷

Actually, this objection appears misled, as we do not require neither a holistic nor an atomistic theory for punctate languages. In fact, we may want to find a way to rule such systems out, for, if we considered them as proper languages, then we would be forced to admit that, in principle, the very notion of language is totally arbitrary. In fact, since we cannot extend the criteria that allow us to talk of infinite languages to finite ones, we would be forced to do the opposite. In this way, every sentence or utterance could count as a language, and this is, as we already noted, definitely undesired. We may choose to have different criteria for natural and formal languages, but thus we would fail to understand on what ground the notion of ‘language’ applies in both cases. Or we can have different criteria for infinite and finite languages, after all. But this is more or less equivalent to rule finite languages out of semantic investigation, for we can stipulate that, since they respond to different criteria, different semantics (if at all) is required too.

We are not interested, then, in finding or arguing for a semantic theory for finite languages. What we require is a semantic theory for infinite languages, i.e., for languages that share the Humboldtian characteristic. Therefore, I will not indulge in much consideration of symbolic codes as human colinguistic gesture, flag signals or semiotically equivalent codes, or animal means of communication. As for the first kind of signalling, it must be noted that there is no right at all to speak of a language, since any structure and any symbolic relation between gesture and world is lacking. All we have is some kind of extralinguistical emphasis added to the linguistic message. If we decide to treat this as a linguistic message, we end up with losing the actual specificity of language.

Something along the line can be said also about the whole lot of semiotic conventions that human beings decide to adopt to interact one with the other without resorting to linguistic communication. Moreover, it is easy to see that in order for any of such conventions to be established and, thus, to work properly, linguistic communication has to be available first.

As for animal communication, it is important to notice what is at stake here. Even if we do not want to pay attention to the fact that animal communication involves very different neurological patterns from human use of language,¹⁸ and that there is nothing in it like the pragmatical features that are essential to the human experience of using a language,¹⁹ we must acknowledge that from an utterly formal point of view there is no point in comparing our language to whatever animals use for their aims. Even the metaphor of animal communication should be understood as merely such—a metaphor. It can be claimed that the reason why acoustic or visual signals among ani-

¹⁷Cf. *ibid.*, pg. 70.

¹⁸Rendall et al. (2009), pg. 235.

¹⁹*Ibid.*

mals function in influencing the behaviour of conspecifics is not their semantic value, but, perhaps, the physical nature of the very sounds and colours.²⁰

I admit that our experience of interpreters can apply to these signals, so that we might look at them as if they were produced within some kind of language and, therefore, as if they were endowed of a semantics. I would not object to such practice, provided that we acknowledge that it might be nothing more than a metaphor deluding us.

The opposite might be true, though. We could be so obsessed with our actual practice of language that we overlook what makes these kinds of signalling similar to our own. Ultimately, since language, whatever our depiction of it, has the characteristic of bearing meanings, semantics must have a role in defining what counts as language. Accordingly, the decision of treating a language relatively to its speaker, or relatively to its lexicon and its rules is only partially independent of the semantic theory we choose.

1.3 Reference, Meaning, Truth

Semantics deals with the way our language is about the world. In this respect we must consider now three crucial notions, viz., those of ‘reference,’ ‘meaning,’ and ‘truth.’ These terms concern different ways our language relates to the world, and a semantic theory is a theory to explain these relations. The term ‘world’ itself can indicate different things depending on which relation we are focusing.

A theory of **reference** is a theory to explain how our *words* relate to *objects* in the world. Under the perspective of a theory of reference, then, we look at the world as made up of objects.

There are semantic theories, like Richard Montague’s, that succeed in correlating every word to an object or an entity of some kind.²¹ Traditionally, these entities will be those described by set theory, such as sets themselves, relations, functions, etc.

In less formal terms, we may say that words have as their references those objects and entities they relate to. One goal of semantics, thus, is – or it might be considered to be – explaining how words relate to their references. To individuate words, anyway, is not so simple a task as we may imagine being accustomed with written languages. Therefore, we might want to talk of *parts of sentences* instead of words.

This move can be useful also in another way, for it exempts us to find a reference for those parts of speech, such as articles or prepositions, which are

²⁰*Ibid.*, pgs. 236 f.

²¹It would be better, for what concerns Montague’s semantics, to talk about ‘pieces of lexicon’ instead of ‘words,’ since this is a theory about formal languages, and the notion of ‘word’ might not be appropriate in this context.

difficultly thought of as having one, especially in a nominalistic perspective. The risk in doing so, on the other hand, is that we might tend to identify parts of sentences by way of their alleged semantic counterparts, thus neglecting their syntactic features, or their syntactic nature altogether.

Another relationship we might want to investigate is that between a *sentence* and some entity we call its **meaning**. We have already seen that meaningful sentences may be said to have *propositions* as their meaning.²² However, propositions are, as noted, highly controversial.

Some try to obviate the problem talking of propositions, hence of meanings, as of *facts* or *events*. What a fact is, though, is disputed too. One standard position is to look at facts as possible combinations of objects, i.e., as states of affairs.²³ Thus, we may call meaningful those sentences whose parts are arranged so as to correspond to a possible combination of the corresponding objects, and meaningless otherwise.

Of course this position requires that we identify references of parts of sentences in the first place. We have already seen this can be a difficult task. It is even impossible, according to those authors who believe that reference is identified by way of sentence meaning. For, if we have first to identify references of the parts of a sentence in order to identify the meaning of the sentence itself, but we believe that references can be identified only after having identified the sentence meaning, we are stuck in a vicious circle.

For what concerns reasons to believe that identification of meaning should come before identification of references, Gottlob Frege wrote:

In the enquiry that follows, I have kept to three fundamental principles:

always to separate sharply the psychological from the logical, the subjective from the objective;

never to ask for the meaning of a word in isolation, but only in the context of a sentence;

never to lose sight of the distinction between concept and object.

[...] If the second principle is not observed, one is almost forced to take as the reference of words mental pictures or acts of the individual mind, and so to offend against the first principle as well.²⁴

²²Cf. § 1.1.

²³Cf. Wittgenstein (1922), § 2.

²⁴Frege (1953b), pg. X. I changed the translation: the word translated with ‘sentence’ is ‘Satz.’ The English translation had ‘proposition,’ which I would not use for a linguistic entity. See § 1.1 for more details.

There are two more risks in wanting to identify facts only after having identified references (i.e., the elements which the fact is made up of):

Seen in retrospect, the failure of correspondence theories of truth based on the notion of fact traces back to a common source: the desire to include in the entity to which a true sentence corresponds not only the objects the sentence is ‘about’ (another idea full of trouble) but also whatever it is the sentence says about them. One well-explored consequence is that it becomes difficult to describe the fact that verifies a sentence except by using the sentence itself. The other consequence is that the relation of correspondence (or ‘picturing’) seems to have direct application to only the simplest sentences (‘Dolores loves Dagmar’).²⁵

Problems about facts and events are well known. We may try to solve them by way of a theory of reference, but we should solve problems in that theory first. Alternatively, we may pursue a theory of meaning by means of a theory of truth. This was, in fact, Davidson’s task. The claim he made in the above quotation is precisely that we cannot arrive to a good theory of truth starting from some theory of facts that looks at facts as compounds corresponding to sentences in every element.

There have always been different views on what **truth** is. It is nowadays accepted virtually by everyone that this notion concerns, rather than the world in itself, the relationship between world and something else – perhaps our thoughts about the world, provided our theories enable us to make such a distinction. More specifically, when dealing with language, we are interested in the truth of those thoughts which are expressed by our sentences. Therefore, a theory of truth has as its goal the explanation of how our *declarative* sentences relate to our *world*.

A discourse is *true* if what it means is actually the case. It is *false* if it is not the case. When we consider isolated sentences, this terminology applies only to assertions, and this is why assertions are the focus of most philosophical research about semantics. Of course, these definitions depend on our theories about meaning. Also, we must be careful in how we identify assertions, since we should not distinguish them from other sentences merely by their capability of being true or false, lest we fall in a vicious circle.

The most effectual characterization of truth has been provided by Alfred Tarski. He defined truth for a formal language \mathcal{L} (called ‘object language’) with respect to another language \mathcal{M} (called ‘metalanguage’), containing a predicate ‘true in \mathcal{L} ’ which cannot be a predicate of \mathcal{L} itself. Let \mathcal{M} contain also a name for every sentence p of \mathcal{L} , which mirrors the structure of p (I will indicate it with the sign $\langle p \rangle$), and, for every sentence p_n of \mathcal{L} , a sentence s_n

²⁵Davidson (2001), pg. 49.

translating it. We can now say, in \mathcal{M} ,

‘ p_n ’ is true in \mathcal{L} if and only if s_n .

This is a very effective definition of truth in artificial languages. It avoids reference to the notion of meaning, because, in the case of artificial languages, we can explicitly construct a correspondence between sentences in the object language and sentences in the metalanguage in a recursive way.²⁶ In a way, Tarski’s treatment of truth comes after an appreciation of meaning; while in the case of formal languages we can surrogate meaning with an explicit correspondence between sentences in the object language and sentences in the metalanguage, such a correspondence cannot be thought of as given in the case of natural languages.²⁷ Tarski’s definition of truth, if applied in the latter case, would involve the notion of meaning, and, therefore, cannot ground it.

Reference, meaning and truth, then, are intertwined notions. It is among the tasks of a good semantic theory to explicit their relationship in order to avoid vicious circles. A theory’s effectiveness in doing so is something we must consider when evaluating it.

²⁶For all this treatment of truth, and for more details, see Tarski (1969), chpt. VIII: “The Concept of Truth in Formalized Languages.”

²⁷On this regard, cf. §§ 2.1.2 and 2.1.3.

Part I

Competitors

CHAPTER 2

Semantic Atomism

Let us now begin our survey of semantic theories with what we will call semantic atomism. It is the theory according to which words or sentences (depending on what we choose to be our atoms) have a significance of their own, independently of context, i.e., of other words' or sentences' significance.

We will see that the choice of atoms is, unsurprisingly enough, somewhat dependent on ontological commitment. Ontological choices on theoretical entities, in their turn, might be influenced by the theory itself, i.e., by its explanatory and heuristic power.¹

The fundamental problem of philosophy of language seems to be: How – and in what sense – can our language be about our world? We may put the question in another form: How can we understand other people's language? It is worth noticing that the two questions, though obviously related, are not the same.

Atomism answers the first question before the second: language is about the world because it consists of a system of independent symbolic relations, and we are able to communicate and to interpret what other people say, in whatever language they say it, because language is about the world, which is by definition one and the same for everyone. Holism first answers the second question, and then the first. We will see how in the second part of the work. For the present moment, let us concentrate on the atomistic answer.

2.1 What Words Mean

Let us begin with an examination of the question we might call the properly semantic one, that about the relationship between language and world. We saw in § 1.3 that this relationship is threefold, depending on which parts of language and of world we focus on. For semantic atomism this is related to our choice of atoms, i.e., to our means of individuation of the smallest parts of

¹Cf. § 1.1.

language that bear the fundamental semantic relation to the world. We must underline that holism and atomism both agree about, for instance, words' having significance. The point in question is whether the semantic properties of words in isolation (or of other atoms, should words prove unsuitable) are fundamental or not.

Semantic holism is a theory according to which none of these relation involving less than the entire system of language can be deemed fundamental: they all depend on the relation between the whole linguistic system and the world. Semantic atomism is the opposite theory: there are fundamental atomic bearers of significance, which ground the relationship between the whole linguistic system and the world (or that part of the world which we can speak about, if such a specification has any sense).

2.1.1 Starting from References

A very classical position about what the bricks of significant language are is to describe the sentence as a connection of terms. For a start, we can think, together with logicians up to 19th century, that both terms and connection are lexically instantiated, respectively by subject and predicate, and by the copula. A semantic theory starting from lexical units, however identified, will have two kinds of questions to answer. The first is about references themselves: what they are and how they are identified (the second part of the question should possibly merge with those about interpretation of language, which constitute, as we have already seen, a different kind of semantic survey). The second kind of semantic questions is about the connection among references operated by connecting the terms referring them, i.e., a theory of the copula.

Words and References

Looking for the basis of language's significance, atomists might want to start assigning it to single portions of sentences, one by one. If the existence of abstract and set-theoretical entities is not problematic for us, we can take the lexical items we have found with the methods described before, and assign a reference to each of them.² Nouns will correspond to objects; linguistic attributions will correspond to properties (perhaps set-theoretically construed); syncategoremata will correspond to set-theoretical functions, operations and relations; and so on. None of these categories contains only concrete entities,

²A word of advice may be needed, lest we think there is some kind of vicious circle involved. Although, in sketching the present kind of semantic theory, I allude to parts of speech such as nouns, no mention of this theoretical term is really needed. We can make the correlation word by word, and then assign all the words referring to the same ontological kind of entities to the same category of parts of speech. Individuation of parts of speech, then, will be deemed a semantic task, not a syntactic or a morphological one. See also the next two notes.

for even objects corresponding to nouns can well be abstract (like ‘whiteness,’ for instance).

It is now time to notice that, although we have reasons to believe we may individuate words in the speech flux, still we might want to keep more vague about what portions of sentences to consider as semantic atoms. As said in § 1.3, in fact, this move would exempt us from assigning a reference to every word in the language. For we might want to avoid resorting to such an ontology as the abstract and set-theoretical one that I mentioned.

We may start with the same connection between words and set-theoretical notions as above. The difference is that now we turn to set-theory just for a methodology, not for an ontology. Once we have the connection we were looking for, we can recursively assign concrete references to parts of sentences. First, take all the words connected with concrete objects (these will be nouns, according to what we said before), and assign them as their reference those very objects. In particular, this will be an easy move to do with proper nouns,³ since their reference is unique. Common nouns, on the other hand, may be thought of as having as their reference some property, which is rather a standard move.

Next, take all the words connected to functions taking concrete objects as their arguments and giving concrete objects as a result (these will be articles, prepositions, some kinds of adjectives, and so on), and compose phrases combining, according to the rules of the syntax, these words with themselves and with the previous ones: the reference of each of these phrases will be the result of the function connected to it.

There will still be problematic words. First of all, adjectives, verbs, and adverbs (words connected with properties, events, and properties of events).⁴ In this respect, let us note that events can, in some way, be traced back to properties of the objects involved in them. This is the reason why Medieval logic was able to recognize in every sentence a subject, a predicate and a copula: even a sentence like ‘Socrates runs’ could be traced back to the allegedly original sentence ‘Socrates is running.’ As for properties, there is no need to look at them as separate from the objects they pertain to.

Next, indexicals, such as pronouns, demonstratives, and some adverbs (‘now,’ ‘then,’ ‘here,’ ‘there,’ and the like). These might be conceived of as derivative from nouns and non-deictic expressions, which they translate and

³When, in this paragraph, I speak of proper and common nouns, I do it with an eye to a naive theory of parts of speech. I have already mentioned (in the preceding note) that we are not at all compelled to such terminology in the kind of theory I am describing (see also the next note).

⁴It is not at all true that, in our languages and in our classification of parts of speech, nouns, adjectives and verbs differentiate because of the kind of entity they are connected to. E.g., there are nouns connected to events (such as ‘war’), or verbs connected to properties (such as, in one of its senses, ‘tower’). The division in parts of speech I am treating here, though, is not the traditional one, for it is based only on semantic grounds.

abbreviate.

Finally, words connected with abstract objects and phrases connected with functions having abstract objects as their result form another kind of problem. There is no easy way to trace abstract entities back to concrete ones, although philosophers have tried to do it for a long time. Alternative ontologies can be resorted to, and this is hardly a semantic issue – though it may have consequences when choosing among semantic theories.⁵

Anything Smaller Than Words?

Before we turn on the semantics of sentences, i.e., of syntactic compositions of words, we should ask ourselves whether any constituents smaller than words can be individuated and exploited as semantic atoms instead of words. Indeed, we know from linguistics that morphemes are defined as the smallest units of speech to which we can attribute some significance. We want to know, then, whether this means that a semantic theory should proceed from morphemes. I think our answer should be no – or, at least, that by doing so no different methodology from that already sketched is needed.

First, let us consider the case of those morphemes coinciding with words (e.g., prepositions). It is plain to see that, to these, the strategy already furnished applies.

In the case of words composed by more morphemes than one, let us assign to each morpheme a corresponding set-theoretical entity, being either a function or an argument, and provide that the final result of the function (or functions) is the entity which was assigned to the whole word by the correspondence of the previous section.

Of course, this way of proceeding has a strong flavour of ad hoc explanation. However, supposing this objection is not fatal, we can appreciate that, as anticipated, the strategy to cope with morphemes is not in principle different from that applying to words. The decision between the two, therefore, is more an ontological than a semantic one: if we are prepared to explain everything in the world by resort to set-theoretical entities, then our semantic theory can start from morphemes. This is so even if we decide, as before, that set theory does not furnish any self-standing ontology, but just a path to the individuation of references. Since resort to morphemes is redundant with respect to words, I will not mention it in what follows.⁶

⁵See below, § 2.3.2.

⁶It is perhaps worth noting, though, that although to choose morphemes as semantic atoms does not bear much significance on the semantic enterprise, there could be some effect on the interpretation of morphological inquiry, or of the morphological features of language. For instance, we may be led into viewing any morphological adjunction as either a modification or a specification of the stem, affecting in some relevant way the connection between the stem itself and its reference. Thus, morphology would be seen more like a semantic matter than, for instance, a syntactic one.

Semantics and Syntax

Once we assign a reference to each lexical unit (or to each of the phrases we want to take as our semantic atoms), we need an explanation of how language assembles words and phrases in sentences so that each sentence's structure results in accord with the intended arrangement of the references of the sentence's words or phrases. In other words, we need an explanation of the semantic value of syntax.

Back in § 1.1 we saw that there are propositions corresponding to sentences, which constitute their meanings. Although a sentence is indeed the kind of entity we would expect to have a structure, we may not think of propositions as structured. But if propositions had no structure, then the structure of our sentences would not help us understand what they mean, and we should bear in mind the infinite associations between possible sentences and meanings. Psychologically implausible as it may be, this scenario is not impossible. However, this means that the connection between lexical units and references does not bear on the connection between sentences and meanings. In other words, the latter semantic relationship must be, in the hypothesis of unstructured propositions, a semantic primitive. This takes us away from the topic of the present section, so I will postpone the discussion of this case to the next one. For the present moment, then, let us assume that propositions are indeed structured, and that we can draw a correspondence between sentential and propositional structure (otherwise, it would be as if propositions were not structured at all).

We have seen how every attribution or predication can be interpreted as having a copula at its origin. If this is the theory we accept, it is easy to see how to derive meanings of sentences. The function of the copula, in fact, is just that of putting together subject and predicate: copulas add nothing to the composition, except for composition itself. The mode of composition (whether it is attribution, predication, relation, etc.) depends on the properties of the lexical units or phrases taking part in the composition, and this contribution can be determined in advance, at the moment of assigning references to words and phrases in isolation.

Theories exploiting the role of copula, anyway, are highly unfavoured nowadays. We introduced the move of tracing every predication to an underlying copula to assign some reference to verbs, adjectives, adverbs and the like, or to phrases having such parts of speech at their core. Of course we may think that this is not at all needed.

We might perhaps renounce to assign references to those elements altogether, thus avoiding theories of subjacency which might prove untenable in the face of empirical data, or even clash with other parts of the theory (e.g., with set theory). We must note, however, that the role of the copula, although referenceless, is not completely void of significance. The copula represents, as already said, the composition of the sentence itself. If all the elements of sen-

tences have a perfectly independent reference, the only way of explaining how they get together to build sentences is to postulate in every sentence the presence of a copula, i.e., of a referenceless element which has the task of keeping the other elements together so that they form the sentence. To get rid of the copula, then, we must find some other element whose reference is responsible for composition, i.e., to keep the sentence together.

In other words, we have two options to explain composition starting from reference. One is to postulate a copula in every sentence, which is responsible for composition: in this way, we can assign to every other element of the sentence but the copula an independent reference. Alternatively, we can avoid appeal to the copula, by letting some elements of the sentence have not independent references.

Frege proposed to charge the predicate itself, treating it always as a function needing an argument, which is furnished by the subject.⁷ Apart from the ontological problem we may find with this proposal, we must underline that we cannot assign a reference to predicates in isolation, since it is only in the context of a sentence that it is showed that their reference must be some unsaturated entity: only in the context of the sentence they show their predicative nature.⁸ In this case, it is perhaps better to renounce definitely to begin our semantic survey with a research of references, and start with sentences' meanings altogether.

2.1.2 Starting from Meanings

Frege's theory of predicates as functions was utterly incompatible with a traditional theory of copula (one which deems the copula to be essential to connect references into meanings, and which, therefore, assumes a copula to be present with this role in every sentence) and with a semantic theory starting from words' references. In general, every theory that derives the semantic role of some element of the sentence from its syntactic role (as Frege derived the references of predicates from their being predicates) must take as its initial field of inquiry at least the whole sentence. Notice this difference between the proceeding described in § 2.1.1 and Frege's proposal: in the former case, we mentioned syntactic categories such as attributions or verbs with the advice that we could avoid such mentions via the ontology of references. On the contrary, Frege's appeal to syntactic roles is unavoidable.

In § 1.3 we quoted a passage by Frege, where he stated his famous *principle of context*:

[...] never to ask for the reference of a word in isolation, but only in the context of a sentence [...].⁹

⁷Frege (1953c), pg. 31.

⁸Cf. *ibid.* and Frege (1953d).

⁹Frege (1953b), pg. X. About the translation, see n. 24 on ch. 1.

Hence, we might decide to be more liberal towards the contribution of the single portions of the sentence and look not for their reference, but only for the meaning of the sentence. All we require now is that sentences' meanings (i.e., propositions) exist and are fixed, and we may take no interest in the references of sentences' structural elements: the real extent of their contribution will be recovered retrospectively.

Let us now take a deeper look into this version of semantic atomism. Our primary interest is, of course, in the connection between sentences and propositions. We have already mentioned that this connection has something to do with the relationship between sentential and propositional structures. The former kind of structure is whatever syntax or logic describe it to be. As for the latter, though, we have already seen that there is no certainty at all that propositions are in fact structured. Let us suppose that they are not. Then, the syntactic or logical structures of sentences are merely devices to help us build them, but have no necessary semantic value: that is, there is no need that a sentence has to be structured in precisely the way it is in order for it to signify the proposition it signifies. Taking this position to the extreme, this means that there might be a language made up of infinite unstructured sentences, each one of them bearing a primitive relation to the propositional world. Although a theory along these lines is highly implausible as a psychological theory, since it gives us no hint at all as for the actual patterns of connection between our language and its meaning, from a formal point of view there is no contradiction whatsoever in entertaining such a view.

There is, though, a problem of economicity. While it is perfectly coherent for a formalistic approach to treat each sentence as isolated for what concerns its meaning, and to consider syntactic and structural analogies and connections among sentences as irrelevant for the semantic enterprise, this is not the shortest way to come up with a semantic theory for all the sentences in our language. For if we consider the sentential structures also as hints to their meanings, then we might trace the meaning of a whole lot of sentences back to the original meaning of only a few of them (in comparison with the original infinity), which we can deem atomic. Such a procedure, though, amounts to consider the meanings of sentences – i.e., propositions – as structured; and, more specifically, we must admit that sentential structure is a means to construe propositional structure.

Thus, let us assume in the following discussion that there are atomic propositions, corresponding to atomic sentences, i.e., to the sentences whose structure is believed to be the simplest possible for a sentence to have. The simplicity I am referring to, though, needs not be a syntactic one: rather, logical simplicity is what we need to let semantics come into the picture. We want atomic sentences to be the logical foundations of meaning. Thus, logically simple sentences will be our atoms.

Our next task is to state which sentences are atomic, and which propositions are their meaning. In their criticism of Davidson's holistic theory, Fodor

and Lepore suggested that good candidates for being considered sentences with an atomic meaning are what they call *token-reflexive* expressions:¹⁰

[...] there is no argument so far that you can't have a language *all* of whose sentences are token-reflexive. [...] If, however, you can have a language that contains only token-reflexive sentences, what argument shows that you can't have a language that contains only one sentence, so long as *it* is token-reflexive?¹¹

I have mentioned a possible reason why we could be suspicious about such a punctate language.¹² Still, Fodor and Lepore's claim that a language might be constituted of token-reflexive sentences, and that the semantics of each of such sentences is atomistic, might not be defied by that argument.

Albeit with a different goal in mind, Quine proposed something similar to explain the first stages of language development.¹³ He proposed a strategy to translate all sentences in our language into a language containing only demonstratives as subjects and their predicates, conceived of in a Fregean style.¹⁴

We need not share Quine's opinion that his strategy has some psychological value. We can just take it as a proof of the logical possibility of reducing all sentences to token-reflexive ones. What happens in human psychology, as we have already said, is another issue. Now, if it is correct that we can trace all other sentences to token-reflexive ones, we can take it for proved that all we need is some atomistic theory for the semantics of token-reflexive sentences.

It might occur to us that token-reflexive expressions were at the heart of one of the semantic theories devised by Bertrand Russell, for whom demonstratives were the logical proper names.¹⁵ Moreover, this kind of sentences reminds us also of the first ideas of logical positivism about language, which stemmed from an interpretation (possibly a misinterpretation, actually)¹⁶ of Wittgenstein's *Tractatus Logico-philosophicus*.

¹⁰See Fodor and Lepore (1992), pg. 224, n. 8:

We're constructing "token-reflexive" broadly to include, for example, tense and any other feature that can function to relativize truth values to contexts of utterance.

Fodor and Lepore do not say where they take the term 'token-reflexive' from. Quine credited REICHENBACH H., *Elements of Symbolic Logic*, New York 1947 (see Quine (1960), pg. 101).

¹¹Fodor and Lepore (1992), pgs. 69 f., emphasis in text.

¹²Cf. § 1.2.

¹³Quine (1960), chpt. III.

¹⁴*Ibid.*, chpt. V.

¹⁵Quine shared this view, to the point that he claimed that "nouns might better have been named pronouns" (Quine (1953), pg. 13).

¹⁶The debate about the right interpretation of Wittgenstein's position as it emerges from the *Tractatus* is endless. However, scholars more or less agree, now, in finding in it much more than merely what positivists would. Wittgenstein himself was very skeptical about Vienna Circle's agenda, and about their consequent interpretation of his work. His own

The way these philosophers explained, at least at one stage of their thoughts, the connection between token-reflexive sentences and facts signified by those sentences was by means of *acquaintance*. This was a sort of logical intuition, a connection between world and language seen as a propriety of a mind, but not of a psychological mind. Some transcendental logical mind was considered responsible for the operations of faculties which showed universality and independence from contingent factors.

Of course, we are not forced to accept token-reflexive sentences as our primitives. Any kind of sentences to which we can reduce all other kinds of sentences will work as our atoms, and we can assume exactly the same (mysterious) mechanism of acquaintance to explain the semantics of the atoms. So, e.g., it is very doubtful that Wittgenstein in his *Tractatus Logico-Philosophicus* had in mind token-reflexive sentences as atoms. In fact, it is at least doubtful whether he conceived his primitive objects (those which atomic sentences are about) as relative to the context of utterance in the relevant way (which they would have been, for instance, if they were sense of data).¹⁷

Such a position makes a strong metaphysical claim about the nature of logic and language. Moreover, there can be some discussion about what is the object of acquaintance. What does it bring us knowledge about? It would be easy for us to answer: the meaning of the sentence, i.e., the proposition that corresponds to it, the facts described by it.

The matter is complicated by the confusion arisen by the neopositivist mention of truth conditions. For it can be demonstrated that all the previous answers are equivalent to the following: By means of an act of acquaintance we know how the world would be if a given sentence were true. Now, it can be proved that, given such a definition, our theory can define only one notion between ‘meaning’ and ‘truth.’ In fact, one is always defined in terms of the other: the Tarskian theory of truth seen in § 1.3 presupposed the notion of meaning of a sentence, and defined truth by putting in correlation a sentence in one language with another sentence in another language. The kind of semantic theory which we are presently surveying, instead, explains truth in terms of some particular fact, described by the sentence, actually taking place. The fact described by the sentence, though, is a logical primitive: it is precisely

opinion about it, however, might not be as definitively relevant as we would deem it. On this regard, cf. this note by G. E. Moore quoted in Rosso (1999), n. 179:

Ramsey had visited [Wittgenstein] in Austria, at the time when he was teaching in an elementary school there, in order to question him as to the meaning of various statements in the *Tractatus*, which Ramsey found difficult to understand. (Ramsey once told me, by the way, that in the case of many of these statements, Wittgenstein said he had forgotten what he meant by them, and had to think hard in order to arrive at a probable opinion as to what he might have meant).

Cf. also Read and Crary (2000).

¹⁷On objects in Wittgenstein’s *Tractatus* see, e.g., Ishiguro (1969) and McGuinness (1981).

what we know by acquaintance.

2.1.3 Starting from Truth

Another kind of semantic theory is possible. It consists in taking truth as a primitive notion, known by acquaintance, and to derive meaning from it. Again, let me stress that the notion of knowledge hinted at when we say that acquaintance is some kind of knowledge has nothing to do with an actual psychological and storable knowledge. It is, instead, a logical assumption we make to explain how language, seen as a logical system of sentences, is connected with the world. Language, according to this as well as to the other theories we have explored in this chapter, needs not be instantiated in human mind as the theory describes it.

The semantic theory we are now surveying should give us truth conditions for every single sentence, independently of every other one. It should be a theory which tells us any sentence's truth value, independently of other sentences' being true or false. Or perhaps we can be more charitable towards semantic atomism: we might require from an atomistic theory of truth that, although the truth values of some sentences are dependent on the truth values of some other sentences, there are atoms, i.e., sentences whose truth values stand on their own.

The derivation of dependent sentences can follow the lines of the Quinean derivation mentioned in the last section. As for the atoms, there are two candidates to this role. First, the truth-conditionally self-standing sentences we are looking for might be the *analytic* sentences of our language. Alternatively, we can think that what we want are the same kind of primitive sentences previously seen, no matter how identified, as long as we have some method to reduce other sentences to them.

These two alternatives, it is better to stress it, are not at all mutually exclusive. Actually, they are probably complementary. Before exploring them with some more attention, let us notice that now we are dealing first of all with declarative sentences, i.e., with that special category of sentences which are susceptible of being true or false.¹⁸ Other kind of sentences (questions, exclamations, etc.) might be deemed derivative for what concerns their meaning. They can in fact be thought of as having the same meaning as the assertions they can be traced back to.

Analytic Sentences

Analytic sentences include two kinds of assertions. First of all, tautologies, i.e., sentences which are always true because they are logically so. More in general, all sentences whose truth depends solely on the language they are formulated in, and not on the way the world is, are analytic.

¹⁸Cf. § 1.3.

Tautologies, being logical truths, rest, in definitive, upon the logical axioms and rules we accept as valid.¹⁹ What constitutes evidence for the validity of logical rules is still disputed. According to the non-psychological perspective we have adopted in the rest of this chapter, we must say that we are not interested in the actual reasons why each one of us accepts or fails to accept them.²⁰ In agreement with the priority granted to truth by the approach we are delineating here, we can consider the truth of logical rules as a primitive knowledge, something known *a priori* by the non-psychological mind.

For what concerns the other kind of analytic sentences (i.e., those sentences whose truth derives from the language they are formulated in), things are not so easy. We are dealing with sentences like “No bachelor is married,” or “All widowers have lost at least one wife.” As we can see, these sentences and the like are founded on the definitions of the terms used in them. Thus, as the truth of tautologies rests on the validity of logical rules, the truth of this larger group of analytic sentences is founded upon the validity of our definitions.

Definitions, in their turn, are a special kind of sentences. In the logical perspective adopted, which is not interested in knowing the historical and social processes that give a word its particular meaning, we might take the action of originally defining as actually having occurred, and consider this particular action the cause of a word having its meaning. The truth of the sentences constituting such original definitions is obvious: since these are indeed at the root of the meaning of at least one term appearing in them, truth cannot be questioned, for there is no way they might be false. As for non-original definitions, they can be thought of as echoing the original ones, thus sharing their truth value: in fact, they are the same sentence. The difference is not in the sentence, but in the illocutive power of the utterance, i.e., in the different goal that makes one utter it. Reasons, though, are not what we are interested in, since reasons are a psychological issue.

Token-reflexive Sentences

We have seen that logical rules and definitions have truth values known by means of what we decided to call an act of acquaintance. Analytic sentences are obviously responsible for the truth of some of our sentences – and, though not exclusively, of all of them. We can argue, in fact, that it is always possible to show that the truth of our sentences agrees with the truth of logical rules and definitions. This is possible because we can always leave inconsistencies to a psychological inquiry, and claim that they are not a linguistic or logical problem. Inasmuch as our sentences’ truth is consistent with logical rules and

¹⁹Given the interchangeability between logical axioms and logical rules, I will refer only to the latter in what follows.

²⁰It is still debated whether some human beings (as individuals or as members of some social group) can show through their reasonings or actions that they consider different logical rules to be valid. Cf. Quine (1960), pgs. 58 f.

definitions, we may deem the latter responsible for the former. In fact, in all cases of inconsistencies between our sentences and logical truths or definitions we will call our sentences false (or appeal to some psychological explanation).

On the other hand, most of our sentences are not analytic, i.e., their truth does not depend solely on logical rules and in the definitions of words used therein. A comparison with the world is quite often needed to determine a sentence's truth value.

In § 2.1.2 we saw how such a comparison is possible, once we know (by acquaintance, as we said there, or by some other means) the meaning of a sentence. We then know, in fact, what to look for in the world to decide its truth value. We also said there that one of the two notions of truth and meaning must be considered primitive with respect to the other, because they are always interdefined. Hence, if we decide to consider truth primitive with respect to meaning, we should find another way to draw such a comparison between sentences and world.

In § 2.1.2 actually there is something else that proved useful for our theory: there we have seen how the semantics of every sentence can be traced back to that of a (possible smaller) sub-group of (primitive) sentences, e.g., token-reflexive sentences, or Russell's $x\phi y$. This claim has nothing to do with the relationship between sentences and world, dealing indeed with the relationship among sentences themselves. Therefore, it is one and the same thing to say that sentences can be traced back to, say, token-reflexive expressions with regard to their meaning (as in § 2.1.2), or with regard to their truth value, which is the issue we are presently interested in.

We can now claim that what we know by acquaintance about token-reflexive sentences is not their meaning, but their truth value. From that knowledge, together with knowledge of logical rules and rules of substitution of words consistent with their definitions, we can know the truth value of every possible sentence.

As for meaning, this is where the connection between truth and meaning comes into play. Let us recall the Tarskian theory given in § 1.3. That theory, as given there, was a definition of truth in a language. In fact, it defines truth thanks to a previously known connection between sentences in a language and sentences in another language, which were supposed to be translations of the first, i.e., to have the same meaning. Meaning, in that theory, was therefore a primitive notion.

We can exploit the same theory the other way round, identifying necessary and sufficient conditions for the truth of a sentence, starting from knowledge of the sentence's being true or false. But atomism allows us to do this only considering one sentence at a time, independently of other sentences' being true or false. Thus, we have on one side a sentence, which we know whether is true or false, and on the other the world, which contains (in some way or another) what makes our sentence be true or false, but we cannot identify what precisely it is. As we cannot be any more specific, we end up with the

quasi-Fregean doctrine according to which all true sentences are true in virtue of the same entity, and all false sentences are false in virtue of the same entity. And, since what makes a sentence true or false is its meaning, this is equivalent to say that all true sentences mean the same, and all false sentences mean the same, viz., Truth and Falsehood themselves.²¹

To recover the idea that sentences with the same truth value might have different meanings, we can make the same move as Frege. He thought that associated to every assertion was not only its meaning (“Bedeutung”, i.e., its truth value) but also its sense (“Sinn”). This notion concerns the way each different sentence is actually true or false. The relationship between truth value and sense, then, is not as strict as a necessary and sufficient condition, as in a Tarskian theory.

Sense is not, according to Frege, a psychological notion, though it is not an entirely logical one either.²² Lacking a more specific way to deal with this notion, we may decide to treat it as a primitive along the lines traced in § 2.1.2, thus nearly equating it with meaning, in our sense of the notion,²³ except for the fact that it cannot be exploited in a Tarskian theory of truth. Or we may think that, after all, it is a quasi-psychological notion; hence, we may rule it out of our present issue, and deal with it when surveying actual use and interpretation of language, in the next section.

2.2 What People Mean

In the preceding section we saw how semantic atomism attempts to answer the first fundamental question of philosophy of language: How – or in what sense – is our language about the world? We turn our attention now to the second question: How can we understand other people’s language? This is a question about language use and language interpretation.

We have often stressed that semantic atomism is not really concerned with psychological issues: this kind of semantic theories borders – and often overlaps with – logical theories. We could deem language an independent logical entity or system, with rules and mechanisms of its own, which we can determine by looking simply at it as a phenomenon, disregarding its origin or purpose. Actually, everything showing the right characteristics could be considered a language and studied consequently: such distinctions as between natural and artificial languages do not matter to this kind of theories.

Use and interpretation of language, though, are a different matter. A survey of these issues, in fact, cannot prescind from language purposes. Anyway,

²¹Frege (1953e), pg. 63. I call quasi-Fregean such a theory because, although the thesis that all true sentences mean the Truth and all false sentences mean the False is indeed Fregean, the argument he constructed to show it had nothing to do with truth conditions and with what a sentence is true or false in virtue of.

²²*Ibid.*, pgs. 35 f.

²³See again § 1.3.

this can mean two very different things: it might imply that we must take into consideration the human nature of language – and either claim or justify the inseparability of language from human mind. Or we can still adopt a more agnostic position, considering language as an instrument connected with some purpose, though not with some deputed users (i.e., human beings). In other words, we might think that language is not necessarily human; therefore, we need not take actual functioning of human mind into account. This last position goes along the lines of the semantic theories seen before.

The question we are addressing here is in its turn twofold, as I briefly mentioned: we might think that we should distinguish between one’s linguistic competence as a speaker and as a listener. In other words, to avoid mention of human roles, we may distinguish between semantic production and semantic interpretation. Semantic production, in a sense, is exactly what we explored in the last section: it concerns, in fact, the way our language produces meaning. We already know at least three major ways to answer this question, depending on what we want to consider as semantic atoms and what semantic notion as primitive.

Let us concentrate on semantic interpretation, then. First of all, let us note that, if we have an atomistic account of semantic production, then the two tasks must necessarily be distinguished. In fact, while production already has the atoms it needs to build up complex expressions, interpretation must start from the complex expression itself. The interpretative task consists in individuating semantic atoms into the complex expression, understanding the way atoms are connected in a significant structure and, thus, interpreting the meaning of the complex expression.

Again, we must consider two different cases. The atomic constituents of the complex expression, in fact, might or might not be accessible to the interpreter.²⁴ I will consider the former case first.

2.2.1 Interpretation by Comparison

If atoms can be used for a comparison with the complex structure to interpret, the process is relatively easy to describe. Of course, to assume that such a comparison is possible, we have to make some idealization about the interpreter, which should be free from limitation of memory, time, and so on. The interpreter can take advantage of different strategies, depending preva-

²⁴By referring in what follows to an *interpreter* it is not my intention to impute the process of interpretation to a human being or to a mind of certain kind. I dub ‘interpreter’ any entity, mechanism, device, or whatever able to interpret, i.e., to recover the significance of expressions.

The process of interpretation here discussed should remain distinct from any psychological process. Whether the mechanisms here described take place in the actual process of human interpretation is a question to be examined afterwards, as it is utterly irrelevant to the present discourse.

lently on the kind of atoms involved and on the kind of semantic primitives (references, meanings, or truth values), as said before.

The easiest case is that, dealt with in § 2.1.1, where syntactic constituents are also semantic atoms. All that the interpreter has to do, in fact, is to compare one by one lexical units in his memory with those in the expression it has to interpret. If syntactic phrases are our semantic atoms, but not lexical units, the interpreter may have to construct those constituents from the lexical units it has identified.

We can assume, now, that the interpreter knows each atom's reference. In fact, if atoms are lexical units, they are finite in number, and it can have in its memory all the connections between atoms and references. If, on the other hand, lexical units are not atoms, the interpreter has nonetheless an effective method to recover references of the relevant phrases (i.e., those phrases which are our atoms) from lexical units, as we pictured in § 2.1.1.

Knowing each atom's reference, and the mode of the connection between them (which is given either by a theory of copula, or by any equivalent theory of composition), it can recover the meaning and, by comparison with the world, the truth value of every expression.

Next, let us examine what happens if we take whole sentences as our atoms. The first thing an interpreter has to do is decompose the complex expression it has to interpret to trace it back to atomic sentences. This can be done only thanks to linguistic means, like lexical hints or individuation of some kind of sentential form, which inform the interpreter that a particular transformation is taking place. The process might involve a certain amount of trial and error. On the other hand, recognition of applications of logical rules or definitions allegedly happens by acquaintance.

The meaning of atoms is allegedly known (by acquaintance or thanks to some kind of Tarskian theory plus a theory of Fregean sense).²⁵ Once decomposition is done, and atoms are recovered, then, the meaning of the complex expression can be built again recomposing it backwards.

Lexical Units

If our atomistic semantics takes words as its atoms, we have another problem to solve: that of individuating the relevant lexical units in the continuity of the discourse. Although all natural languages are passible of being written, their written form presupposes lexical segmentation (i.e., breaking up speech into lexical units), and cannot therefore be used to support it. The obvious field to look at for answers regarding what should count as a word (a lexical unit) and what should count either as part of a word or as a compound is, nowadays, linguistics.

²⁵See, respectively, § 2.1.2 and § 2.1.3.

It cannot be denied that semantics has an effective role in actual identification of lexical items. This route of investigation, though, is precluded to us, because it would require us to already have the semantic theory we are attempting to found. Fortunately, there is a sense in which actual processes might not be what we should be concerned with. Everyday practice of language could in fact obliterate some true fact about language itself. For natural languages depend on contingencies such as human biology and human history, which could well affect our daily usage of language, but which we can deem scarcely of any importance regarding what language really is—regarding its essence, so to say. Provided we can find an alternative explanation, what happens in everyday interpretation of language may be labeled as a psychological problem.²⁶

The most useful answers about lexical segmentation come from phonology. Instrumental measures assure us that there is no audible pause between words in fluent speech.²⁷ Nonetheless, there is some evidence²⁸ that phonotactics (i.e., knowledge about frequency of patterns of phonemes)²⁹ and metrical information (i.e., knowledge about rhythm and stress pattern), together with data from pauses bounding utterances (which are audible), furnish some good information to itemize speech into words.

Such information, of course, is not the only available one, when we have to individuate words in speech. We can count also upon our preexistent knowledge of what is a word. But this is not so for infants, and this is not so for artificial computing systems. Both infants and artificial systems, though, can get quite good at individuating word boundaries, and, at least when the task is accomplished by artificial devices, we know that no other information rather than phonological one is required.³⁰ This is all we need to claim that the task is generally speaking accomplishable.

Linguistic data are indeed holistic: to interpret them as relevant data, the linguist has to contrast each single datum against the context of (possibly all) other data, and only from the evaluation of all the linguistic system data show their importance. This has been true ever since modern linguistics moved its first steps.³¹ In particular, this must be the case with the kind of phonological evidence that we are invoking here. However, this holistic nature of linguistics

²⁶Of course, there are theories for which exactly these elements that we have just dubbed of scarce importance (biology, history, psychology) constitute the essence of language. These theories will have problems in addressing semantics from the point of view that is now being described. We may call the present perspective the logicistic one. For further comments on the strong relation between semantic atomism and a logical perspective, see § 2.3.2.

²⁷Cairns et al. (1994), pg. 32.

²⁸See, e.g., Mehler et al. (1981), Cutler and Norris (1988), Cutler and Butterfield (1991), Cutler and Butterfield (1992), Rao (1993), Gow Jr. and Gordon (1995), Gow Jr. et al. (1996), Davis (2003).

²⁹But see Cairns et al. (1994).

³⁰See Rao (1993).

³¹Cf. Morpurgo Davies (1996).

is uncommitting for semantic atomists, since a holistic point of view about phonological theories does not entail semantic holism.

Phonology, as all human science, is not infallible. Neither its data are always as univocal as philosophers might want them to. However, with some good scientific theory working on the issue, semantic atomists can well assume that words can, at least in principle, be individuated without semantic knowledge. Notice that if this were not the case, and if semantics were indeed necessary to lexical individuation, then we could not have an atomistic theory of interpretation. In fact, we would need to individuate words in order to know what they mean and what the sentence they appear in means, but we could not individuate words unless we knew what they mean.

2.2.2 Interpretation without Clue

Let us assume, now, that the interpreter cannot compare atoms available to it with the complex structure to interpret, in order to find out which atoms it is made up of.³² I talk of generical availability for two reasons. First, I do not want to make any assumption about the nature of the interpreter: so I keep vague about the mechanisms of this availability. Secondly, different linguistic constituents can be deemed to be semantic atoms, and they are available in different ways:³³ so, it is better to keep vague about such a notion of availability.

If atoms are not available, they must somehow be recovered, and this can be done by making two distinct ways interact: one is decomposition of the complex expression in its syntactical constituents, the other is looking at the world. The individuation of constituents could seem a simple syntactic matter. Indeed, some facts about word order, or movement of phrases from one point of the sentence to another, or agreement, etc. are quite independent from semantics. Nevertheless, although the composition of phrases starting from lexical units can be defined a matter of pure syntax (at least according to some syntactic theory, like the classical generative one), individuation of such constituents (i.e., phrases) in the sentence cannot. Thus, production and comprehension of language must be treated as two quite unrelated tasks. For, while the individuation of words in the complex expression is (or can be construed as) utterly independent of semantics (as seen in § 2.1.1), this is not so for what concerns the assignment of words to different parts of speech. In § 2.1.1 we were allowed to make arbitrarily such an assignment, because

³²This can happen for a variety of reasons: for instance, the interpreter may lack lexical memory, or the expression may be constituted of atoms that differ from those available to the interpreter.

³³In fact, while a finite number of atomic constituents might always be available to some kind of memory, an infinite number of them should be syntactically built, thus being only potentially available. Of course, in the latter case, syntactic constituents would not be deemed semantic atoms.

we assumed to be starting from an intuitive knowledge of references. This is precisely what we have chosen not to assume here.

Therefore, interaction with world knowledge and world survey is needed. The interpreter must be able to find in propositions (seen as real entities of the world) elements corresponding to the elements of sentences – i.e., their references. Only when it has devised a method to assign each word to a part of speech, it can construe phrases and complex expressions. This, however, can be done only through a trial and error procedure, and there is no guarantee that a unique method of doing so exists. In fact, the contrary can be proved.³⁴

Individuation of other kinds of semantic atoms (i.e., token-reflexive sentences and analytic assertions) is perhaps possible. This implies that the interpreter manages to recover such atoms in the complex sentence. Since it cannot do it by direct comparison, an indirect one might work. A possible way, then, is to know atoms of a language different from that in which the expression to interpret is formulated. Then, by trial and error and some behavioristic survey, the interpreter might manage to build a translation theory and to interpret the expression.

2.3 Evaluating Semantic Atomism

Up to this point we have seen how atomistic semantic theories manage to answer questions about the relationship between language and world. It is now time to look if those answers are in fact valuable and profitable ones, and what is their role in the landscape of the different modern disciplines dealing with one aspect or the other of language.

2.3.1 Internal Problems

I dub ‘internal’ those problems that arise inside a theory, before any interaction between the theory itself and any other theory or any other discipline. Contradictions, incoherences, vagueness, etc. are all instances of what I mean with the expression, and these I will now examine.

Reference-grounded Theories

The major problem with reference-grounded theories is that the compositionality attributed to the language is too simple to be a good explanation of natural semantics. Our intuition tells us that the same word assumes different meanings and even different semantic values in different contexts, which is Frege’s objection to any correspondentist semantics grounded on references. If we want to follow this intuition in an atomistic theory, we end up with having to adopt one of two solutions, both of which appear to me quite unsatisfactory.

³⁴Cf. Quine (1960).

We can split different occurrences of the word among what we now are forced to interpret as different atoms, each with its different referential relation.³⁵ We would thus lose trace of the common ground which allowed us to speak of a single word in the first place. Of course, there is no need that a theory like the one we are devising should account for something that, after all, is only in intuition and common practice. Theoretical semantics can very well override intuitive semantics. Still, we need a principled way to know when to stop finding different references.

The reason I think there is a problem will be apparent after the examination of the other solution I mentioned before. In a sense, it is just a specification of the premises hidden in the previous solution. We could, that is, assign as a reference to any intuitively ambiguous term a function which takes as an argument the reference of the terms that enable us to solve the ambiguity. This is, I think, the core of Frege's proposal, which consists in putting the context inside the reference.

I have two objections to such a methodology. First of all, I doubt that we are still entitled to talk of atomism, here. Putting the context inside the reference is just a roundabout way of taking the context into consideration when assigning a reference to a word, a clumsy way not to admit that a word's reference is dependent on at least some portion of the context.³⁶ And even if we try not to explicit this role of context, and stick to the first kind of solution proposed, there is still a role of the context that we cannot amend: for context has at least a heuristical import on our enterprise. If we did not take context into consideration we would never be led to consider what appeared as the same word as, in fact, different words with different referential connections to the world. But if this is true, then we must admit that an atomistic theory for natural languages is only the axiomatic assessment of semantic relations that we find through other ways. In other words, there is no insight to the nature of natural semantics in this kind of atomistic theories.

Meaning-grounded Theories

Of course, the most problematic notion involved in the second kind of semantic theories we have examined is that of acquaintance, which is a rather mysterious interaction between a subject of knowledge and what it comes to know. This, however, is not an internal problem of the theory: a metaphysical theory to support such a notion can in fact be devised, and, as long as this possibility exists, semanticists as such need not worry.

³⁵For the sake of simplification, I will speak here of semantic atoms as if we already decided that words or lexical items should count as ones. I see no easy solution coming to the problem from the adoption of syntactic complexes as semantic atoms.

³⁶The possibility that the relevant portion of context is the sentence, as per § 2.1.2, is still open, of course.

But other difficulties exist. We have seen in § 2.1.2 that it is possible to trace every sentence of our language back to those having demonstratives as subjects. However, the mechanism of this reduction is not always as clear and flawless as we might desire. Idiomatic sentences are the less problematic ones, although their reduction cannot happen in an utterly mechanic way. More problems come from the treatment of metaphorical and, in general, figurative speech, where sentences and phrases exploit, so it seems, their face meaning to convey some meaning of second degree.

The greatest problems, however, come from **intension**. By this term we mean the characteristic shared by all those kinds of sentences for which the following rule does not hold: in a sentence, substitution of a phrase with another one referring to the same as the first does not change the meaning of the whole. For instance, given that ‘Cicero’ and ‘Tully’ refer to the same person, the two sentences ‘Tully wrote *Pro Archia*’ and ‘Cicero wrote *Pro Archia*’ have the same meaning; the two sentences ‘Mark believes Cicero to have written *Pro Archia*’ and ‘Mark believes Tully to have written *Pro Archia*,’ however, do not have the same meaning, if we take the two sentences as expressing the content of a belief of Mark’s about some fact, and not about a precise individual. The same happens with modal sentences: given the identity ‘the number of planets of the Solar System = 8,’ the two sentences ‘It is necessary that $3+5 = 8$ ’ and ‘It is necessary that $3+5 =$ the number of planets of the Solar System’ have different meanings. Modalities and propositional attitudes (i.e., attitudes towards propositions, like believing that, or thinking that, or fearing that, etc.) always give rise to *intensional contexts*.

The problem with intension is that it forbids reduction of embedded sentences: after all, there is a difference in meaning between ‘Mark thinks that the ball is blue’ and ‘Mark thinks that this is a ball and this is blue,’ although ‘The ball is blue’ allegedly reduces to ‘This is a ball and this is blue.’ Reduction of intensional sentences, besides, is very difficult. To reduce intensional sentences to *extensional* ones (i.e., a sentence for which the above-mentioned rule holds), one has to swallow ontological entities such as possible worlds, sets of propositions, etc.

There is of course no contradiction in entertaining such an ontology, although one might prefer not to. But the problem with this kind of solution is, seemingly to what we saw about reference-grounded theories, that, as we saw in the case of ‘Mark believes Cicero to have written *Pro Archia*’ and ‘Mark believes Tully to have written *Pro Archia*,’ there is an ambiguity that makes some sentences express the same meaning or not, depending on the way we interpret it. The possibility of different interpretation must be written in the semantics of the sentences. There is more than an ontological problem, then. If we want to account for intensional sentences in an extensional way, we must somehow redeem natural languages from their ambiguities. The alternative would be, thus resolving also the problem concerning idiomatic expressions and figurative speech, to reserve this kind of semantic theory to some special

kind of language, leaving natural languages with all their complications to some different theory, perhaps a theory of interpretation.

Truth-grounded Theories

Subdivision of declarative sentences into analytic and non-analytic (i.e., *synthetic*) ones might seem *prima facie* obvious. In fact, it is far from being so:

The word ‘postulate’ is significant only relative to an act of inquiry; we apply the word to a set of statements just in so far as we happen, for the year or the moment, to be thinking of those statements in relation to the statements which can be reached from them by some set of transformations to which we have seen fit to direct our attention.³⁷

The analyticity of one sentence, then, depends on the relationship between that sentence and other ones (or in what we consider this relationship to be). It cannot be explained, then, by considering the analytic sentence in isolation. I take this to mean that, just like in the cases of reference and meaning, our only guide to establish the truth of sentences is our intuition of it and, perhaps, formal logic, which is, anyway, a guide only for a limited class of assertions, viz., tautologies.

We might think of solving this problem freezing language in one single stage: although we know that there is no way to make a distinction between analytic and synthetic sentences, we can decide to draw an arbitrary and artificial one. After having done so, we can afford reduction. Of course, there will be problems connected to idiomatic and intensional sentences, as we have seen. Nonetheless, once we have put ourselves in an artificial situation, there is no obstacle to artificially solving those problems too, in a way similar to what we have already seen.

The major problem of such a theory, anyway, is that it is not at all clear how to go from knowledge of the truth value of a sentence to its meaning. We have seen that the Tarskian theory is of no help, unless we postulate some Fregean sense of the sentence distinct from both its meaning and its truth value, but somehow connected with both.

Besides, this kind of theories assumes that all sentences can be traced back to assertions. While this is a relatively easy job for questions, other kinds of sentences are more complex. Promises, or ironic sentences, and in general the vast area of figurative speech need some more detailed explication.

³⁷Quine (1953), pg. 35.

Theories of Interpretation

Let us now focus our attention upon the group of theories we have examined in § 2.2. Notice that these are not properly semantic theories. Let me quote a passage from Fodor and Lepore's book on semantic holism, to show what the relationship between semantic theories and theories of interpretation is likely to be:

Meaning (or semantic) holism must be distinguished from a number of related ideas with which it's easily confused: from holism about confirmation, for example, or about interpretation or about the individuation of functional properties. These other kinds of holism might be true even if meaning holism isn't. Or at least, so it seems; unless there are arguments to show not just that confirmation, interpretation, or the individuation of functional properties is holistic but also that the meaning of a symbol is somehow determined by facts about how its applications are confirmed or how its tokens are interpreted or the functions that it performs.³⁸

In § 2.2, we have been forced from time to time to resort to holistic or quasi-holistic means. Any trial and error comparison, in fact, presuppose that the whole set of what it is to be compared and trialed is available. Therefore, recognizing the significance of any expression is always, for any interpreter, a partly holistic procedure. Significance itself, however, is not necessarily holistic: for a sentence to have some meaning, or a word or phrase to have a reference, there is no need that an entire system of meanings and references is assumed. To interpret the same sentence or word or phrase such an assumption, so it seems, must be done: therefore, theories of interpretation might be very different from theories of meaning.

2.3.2 Interactions

We turn now to the place of semantic atomism with respect to theories of other disciplines. This is something important in evaluating a theory, because it gives us a relative measure of its value and defensibility. Although this is not a universal value, we may choose to improve those theories that are closer to certain areas of research and not others.

Semantic atomism seems to be very close to some metaphysical and logical issue. The questions it attempts to answer are addressed also by linguistics and psychology. Thus, we will try to delineate the relationships between semantic atomism and these four disciplines.

³⁸Fodor and Lepore (1992), pg. x.

Very strong metaphysical and ontological commitments are needed to sustain an atomistic position towards semantics. We have not failed to stress things being so in the discussion above.

The ontology of reference-grounded theories, as we have seen, is a very rich and complex one, assuming some kind of existence for abstract entities and set-theoretical ones, such as concepts, classes, functions, and so on. We may have various reasons for not believing in some kind of existence of these: we may fail to assign them certain identification criteria, for instance, or we may think that mere usefulness in a theory cannot be a proof of existence.

In fact, we have seen that there are some ways we can overcome the problem: we can consider the most problematic entities, such as set-theoretical ones, merely for proof's sake, and claim existence only for the unproblematic ones. This might seem an easy move to do, but in fact it must cope with the problem of properties. A theory excluding set-theoretical and abstract entities as much as possible, cannot consider properties as distinct from the entity the property is a property of. This, however, leaves us with some metaphysical problem of not simple solution. The way we identify what standard metaphysics dubs **substances**, i.e., substrates of properties, is usually by means of nouns, or at least some nouns, like proper nouns or nouns of natural kinds (e.g., 'water,' 'gold,' 'oak tree,' and so on). However, it can be claimed that the word 'gold' does not refer to a substance, but to a property, the property of being gold; and the same for all nouns, even for proper nouns. Thus, it seems that the only way to refer to substances is by means of indexicals ('this' or 'that').

Two problems arise. The first is that it seems doubtful that the references of indexicals are actually substances: for indexicals can perhaps be construed as owing their reference to spatiotemporal properties, such as being in a certain place at a certain time. And if this is correct, we must admit that the percipient subject is the only substance to which all properties are relative.

If this is true, then, and that is the second problem, references can only be identified through complex phrases. Eventually, if properties are not independent from substances, each sentence may consist in a single referring expression. There is no difference, then, from considering references or meanings as primitives. In other words, talking of references may be merely redundant.

But reference is not the only notion that can arise ontological problems. Even sticking to meanings, we may commit to ontologies involving rather objectionable entities (possible worlds, functions and other set-theoretical entities, propositions, abstract properties, and so on); not to count the notion of a non-psychological mind having acquaintances and knowledge. We might need a feasible metaphysic to deal with these notions and to give them plausibility.

Under a different point of view, though, such a metaphysical theory may not be necessary. If we consider the even stronger relationship between atomistic theories of meaning and logic, metaphysical and ontological assumptions might

be regarded as merely methodologic. No further commitment to them would be meant; they would not bind us. In other words, we may assume a rich and plentiful ontology just to support a theory which can prove of some utility in logical matters, while we can still be more austere as for what we believe the universe actually consists of.

The strongest relation entertained by semantic atomism, thus, is with logic. We have seen, indeed, the strong anti-psychologicist claims moved by the theory. This makes semantic atomism an easy prerequisite to be met for logical theories. These, in their turn, can largely improve the semantic theory by furnishing it with the rules of transformation and reduction it needs.

The side effect of this harmonious relationship between atomistic semantics and logic is that, as we have already noted, semantic atomism is not quite good at describing semantics of natural languages. As a consequence, while semantic atomism needs some portion of modern linguistic theories to get developed (such as phonology or syntax or maybe even morphology), it has very little to offer in exchange. Modern linguistics is a deeply holistic discipline, in every aspect and portion; therefore, an atomistic theory cannot help linguists solving their problems, nor furnishing them with an agenda.

As for psychology, it is clear that, as long as semantic atomism carries on with its claims that all actual processes involving language are a psychological issue and not a semantic one, the two disciplines cannot cooperate. Nevertheless, what we have discussed in § 2.2 is a quasi-psychological side of the matter. The mention of behavior in the last paragraph of that section is emblematic. Although it does not purport any assumption about the nature of the interpreter or the producer of linguistic expressions (for mechanisms and devices have a behavior too), still it gives to all the discussion a naturalistic flavour, missing in the properly logical part (i.e., theories examined in § 2.1). Besides, although this mention comes only in the last paragraph, it might appear clear that all the section describes the process of interpretation from a behavioristic point of view, and the very claim of not making assumption about the nature of the interpreter is completely behavioristic in tone.

2.4 Some Brief Conclusion

Semantic atomism is a theory about language, not about human language, nor about humans using language. Every assumption about the nature of language that comes from language being a human feature must be neglected by any atomistic semantics.

Atomistic theories are far from being complete or entirely convincing. Their flaws, though, are worthy of mending, because of the strong association between semantic atomism and modern logic.

Ontological and metaphysical weaknesses of such theories can be ignored in

the light of the logical value of them. Nevertheless, there are clear indications for metaphysics that some work still needs to be done, and that an answer to metaphysical questions about universals, abstracts, and knowledge can prove useful.

Still, semantic atomism has little to do with natural languages and, therefore, with linguistics. Moreover, it cannot furnish us with a gnoseology of language, nor with a theory of linguistic interpretation, without giving up to holistic elements.

CHAPTER 3

Semantic Molecularism

The second approach to semantics that I would like to sketch is the molecularistic approach. While atomists claim that meaning can be accounted for by paying no attention to the way languages constitute a system, molecularists hold a somewhat opposite view: they think that to understand the meaning of any atom of meaning (whatever partition of language we choose to define as atomic) we need to consider the meaning of at least a part of the other atoms in the language.

This point of view springs from considerations about the peculiar character of the semantic enterprise, and about the connection between the two parts of philosophy of language, as we previously characterized them. Again, this point is rendered clear by a comparison with atomism. In the previous chapter we saw that a possible semantic theory (more properly: a possible set of theories) consists in distinguishing between a proper semantic answer to the abstract problem of meaning and some other kind of theory, perhaps bordering philosophy of mind and psychology, about the way we manage to attach such an abstract meaning to our language: that is, a theory of interpretation. We have reached the conclusion that such an approach, while not in itself mistaken nor inconsistent, falls short of the desired goal of giving both a theory of meaning and a theory of interpretation. The first move towards a molecularistic semantics is to acknowledge that the two theories, though still somewhat distinct, are so intertwined that it makes little sense to face them one at a time:

A step in the right direction is taken by conforming to what is, I believe, the correct observation that philosophical questions about meaning are best interpreted as questions about understanding: a dictum about what the meaning of an expression consists in must be construed as a thesis about what it is to *know* its meaning.¹

¹Dummett (1976), pg. 69, emphasis in text.

From this passage by Michael Dummett it is very clear the interdependence of the two order of phenomena we have delineated. Indeed, it seems that Dummett says more, arriving to claim that they are the same kind of phenomena, and that it is a misconstrual of the semantic enterprise to assign different scopes to a theory of meaning in the narrow sense and to a theory of interpretation. This would go against the outline I tried to give to what I perceive as the core problem of meaning in the philosophy of language. But even if Dummett says that the whole process of establishing a theory of meaning should be construed as a foundation of a theory of interpretation, he himself gives us a twofold theory, which accounts somewhat separately for the meaning and for the use of sentences. The heuristic key to the foundation of the properly semantic part of the theory, though, comes from a reflection on how we recognize the meaning of sentences.

3.1 The Meaning and Sense of Sentences

From what Dummett writes, it may seem that the reason he believes his theory to be a molecularistic one is the stress he puts upon sentences instead of words. He thinks, that is, that the real semantic atoms of the language are words, but that words have meaning only when used in sentences.² This is no surprise, given the Fregean background of Dummett's philosophy. We might think, though, that to start from sentences is not a sufficient reason for seeing a theory as molecularistic. After all, sentences can be seen as bearers of a primary semantic value, from which words' semantic value can be construed as merely derivative. This was the core of the approach examined in § 2.1.2. Dummett's theory is nonetheless molecularistic, since other features in it drive us to attaching such a label to it.

The most important of these features is what we can call the distinction between the **meaning** and the **sense** of a sentence, with a terminology which is Fregean, although maybe a little stretched, since I will use the terms somewhat differently. I would say that the *meaning* of the sentence is the properly semantic notion, and it is abstractly identifiable independently of the sentence's connections with other sentences. Such an identification needs not be an effective identification; we might very well not have any idea of the meaning of a sentence in isolation. The fact that we may still be able to consider such a notion is enough to consign it to our inquiry.

What I call the *sense* of a sentence is, on the other hand, identifiable with the way the sentence is used, with that aspect of the world and of human practice that is related in a peculiar way to the sentence itself, and the recognition of which may prove relevant for the understanding of its meaning. Such a notion is to be distinguished, in its turn, from that of **force** of a sentence, which is the most concrete of the three, for it concerns the actual usage of lan-

²*Ibid.*, pgs. 71 f.

guage, independently of its relevance for the attribution of the understanding of meaning.

The relation which connects the notion of sense to that of meaning is the relation of grounding. Michael Dummett takes into consideration three different basic notions which might function as a founding ground for a molecularistic semantics. These are truth, verification, and falsification. Although he manifests some preference for the last one, he claims his discourse to be not completely conclusive in this direction. Truth is, the other ways entail difficulties that, maybe, are not faced when dealing with falsification. But let us start with truth nonetheless, for even if we might be forced to abandon that path, it will still prove of some utility to have walked along it for some length.

3.1.1 Meaning As Truth Conditions

The traditional semantic theory which Dummett considers first is the one which starts from the identification of the meaning of a sentence with its truth conditions, i.e., with the state of affairs which would render an utterance of the sentence true. Up to this point, there is no difference with the atomistic theory sketched in § 2.1.3. There, we faced the problem of defining the way such a state of affair could be identified, and our answer appealed to some sort of primitive acquaintance.

Dummett is unsatisfied by such an answer, and one of his points is that the desire to ground the semantic enterprise on the concept of truth faces its difficulties precisely in the grounding moment:

Thanks to the work of Frege, Tarski and many others, the difficulties that face the construction of a theory of truth conditions are not difficulties of *detail*: they are difficulties of *principle*, that face us at the very outset of the enterprise. We know well enough how to build the machine: but we have no grounds for confidence that we can set it in operation. There are some problems of detail, of course, concerning the adaptation to natural language of the techniques devised by Frege and Tarski for formalized languages; but we may reasonably feel optimistic about finding solutions to problems of this kind.³

In the answer Dummett tentatively gives to the foundation problem lies the molecularism of his position. In fact, he does not think that to solve the problem of assigning a reference to the words in a sentence we only have to consider the whole sentence in which they figure, but that we must take into consideration also a whole possibly infinite set of sentences where those words appear or might appear. The key to such a conclusion is in the formulation

³ *Ibid.*, p. 68, emphasis in text.

of the problem of meaning in terms of interpretation. What we are interested in is not the meaning of an abstract sentence, but the construal of a specific act of uttering it, because only the ability of uttering it can be a display of the right knowledge of the meaning of the abstract sentence.

Someone who knows, of a given sentence, what condition must obtain for it to be true does not yet know all that he needs to know in order to grasp the significance of an utterance of that sentence. If we suppose that he does, we are surreptitiously attributing to him an understanding of the way in which the truth condition of a sentence determines the conventional significance of an utterance of it: but, since the theory of meaning is intended to display explicitly all that a speaker must implicitly know in order to be able to speak the language, the presumed connection between the truth condition of a sentence and the character of the linguistic act effected by uttering it must be made explicit in the theory.⁴

In other words, the difference between the atomistic theories we examined in the past chapter and the theories we are taking into consideration presently (together with those which we will be dealing with under the label of semantic holism) is that in the former case we began by dealing with abstract metaphysical entities deemed responsible of the semantic features of the language. Thus, we have said how the major obstacle to entertaining an atomistic semantic theory is not its inner inconsistency, but its uselessness with regard to the interpretative task. The key to develop a theory that helps in this respect is to begin with interpretation, or, at least, to begin not with abstract entities but with observable phenomena. That is, sentences can still be taken into consideration, but their semantic features must be deduced from their displays, i.e., from utterances.

Although we still have to characterize a form of holism that appears feasible, we can already ask what is the difference between such an approach and the molecularistic one which we are trying to sketch in this chapter. After all, even without further details, we can anticipate that holism too sets as a starting point for its enquiry on semantics the comparison of all the possible uses of utterances, as a means to construe the meaning of sentences. The core difference is that holism cannot distinguish among different sentences before having taken all of the utterances into consideration.⁵ Molecularism, on the other hand, is somehow two-headed, because from one side it identifies sentences, and from the other it interprets such sentences' meaning by surrounding each sentence with all the possible utterings of it.

⁴*Ibid.*, pg. 73.

⁵We will have to see how this task is possible, since we have already showed that it is at least disputable whether such a thing as taking *all* of the utterances into consideration is possible at all.

There is also another, more traditional, difference between the holistic and the molecularistic views:

The difference between a molecular and a holistic view of language is not that, on a molecular view, each sentence could, in principle, be understood in isolation, but that, on a holistic view, it is impossible fully to understand any sentence without knowing the entire language, whereas, on a molecular view, there is, for each sentence, a determinate fragment of the language a knowledge of which will suffice for a complete understanding of that sentence. Such a conception allows for the arrangement of sentences and expressions of the language in a partial ordering, according as the understanding of one expression is or is not dependent upon the prior understanding of another.⁶

The reason why this difference is, in my opinion, only accessory and should not constitute the core difference between a molecularistic and a holistic theory is that it does not help distinguishing molecularism from atomism. What Dummett writes about molecularism, in fact, can be said also about any serious semantic atomism. Few atomists would maintain, that is, that every single sentence is utterly independent, for what concerns meaning, from the others. But they would claim the existence of atoms, which are, in Dummett's analogy, the minimal elements of the partial ordering, whose interpretation does not call for further sentences or expression.

I do not want to suggest that Dummett was misled about his own theory: he had the right to define molecularism by any feature he found useful. I only want to say that, even if I refuse the definitions he furnished and I try to provide my own, I still find his theory a molecularistic one, perhaps with some adjustment that, by now, should be clear. For instance, I do not think that in a molecularistic theory there is any point in speaking of partial ordering and, hence, of minimal elements.⁷ I think, though, that the traditional difference between molecularism and holism points out an important feature of the former, i.e., the possibility of dividing the whole language in several sub-systems of sentences which do not call for sentences from the outside as for determining meaning. Whether this feature is defensible or even necessary to any semantic molecularism is a question I will face later on.

⁶ *Ibid.*, pg. 79.

⁷ Here there is a problem connected not only with interpretation but also with the acquisition of language. Language, in fact, is not learnt all of a sudden, but gradually; what is more, the part of language that we have already learnt has a role in the acquisition of other subsequent parts. In this sense, with regard to this priority of learning, we may be tempted of finding those minimal parts I am arguing against.

I do not think that this kind of objection is unsurmountable, but to deal with it presently would take us away from the issue of this section. Therefore, I postpone the matter until § 6.2.2. Although there I consider the matter from a holistic point of view, I think that it would be easy to rebuild that argument in a molecularistic fashion.

3.1.2 Meaning As Verification

In the previous section, we considered truth condition to be identified thanks to a whole complex of sentences surrounding the one we are focused on in a given moment, and we said that the construal of such a complex of sentences depends in its turn on utterances (i.e., on the use we make of the sentences themselves). Now it is time to mention that we can indeed doubt that what we obtain through such an interpretive procedure are truth conditions. For we understand truth conditions to be independent from our interpretation of sentences expressing them; therefore, we think of them as independent from our ability of recognizing true sentences at all.

Dummett acknowledges the problem, which arises clearly in presence of sentences in the future tense, or with conterfactuals, or again with sentences concerning some kind of infinity. In all such cases, the possibility of assuming a controversial metaphysical stand for realism is certainly open:

This line of thought is related to a [...] regulative principle governing the notion of truth: If a statement is true, it must be in principle possible to know that it is true. [...] [F]or, if it were in principle impossible to know the truth of some true statement, how could there be anything which *made* that statement true?⁸

From such a position, though, we are exposed to any objection which asked us to point to the facts which make a given sentence true. However, we may try to put the stress on the recognition of truth, rather than on truth itself. Thus, we notice that the relevant notion is not that of truth condition, but that of verification. Assuming such a notion as central for our inquiry must lead us to abandon the pretense to apply classical logic to the semantics of natural languages. For in the cases contemplated above, of undecidable sentences, the principle of bivalence does not hold: in other words, there is no way to maintain that such sentences have a determinate truth value out of two (or more, in the case of multi-value logic).⁹ Realism about truth condition is a metaphysical position which, so far, we cannot support with any proofs.

When we begin to think that the relevant notion is that of being capable of effectively assigning a sentence a truth value, we realize that such a task cannot be accomplished for all sentences. Since our intuitions about natural languages are perhaps different (i.e., since we may be lead to think that every statement we make is either true or false), Dummett concludes that the semantic theory we are trying to found will require a revision of our linguistic practice.¹⁰ We may draw a somewhat weaker conclusion, claiming that it is only our intuitions about linguistic practice that need revising, or perhaps our understanding of what is intuitive to claim about our linguistic practice.

⁸*Ibid.*, pg. 99, emphasis in text.

⁹*Ibid.*, pg. 102.

¹⁰*Ibid.*, pg. 103.

To illustrate what kind of semantics he has in mind, Dummett calls our attention on intuitionistic explanation of the meaning of mathematical statements:

In this case, there is no problem about the meanings of atomic statements, namely, in this context, numerical equations, since these are decidable: a grasp of their meaning may be taken as consisting in a knowledge of the computation procedure which decides their truth or falsity. The whole difference between the classical or platonistic and the intuitionistic interpretation of arithmetical statements therefore turns upon the way in which we are given the meanings of the logical constants—the sentential operators and the quantifiers.¹¹

The explanation of such difference wipes the ambiguity arisen by the mention of “atomic statements” out. Here Dummett is perhaps hinting at the role of such statements in the arithmetic theory, or perhaps even in a semantic theory given in a classical or platonistic way. Surely, there is no such thing as atomic statement in the sense of an atomistic semantics in the kind of intuitionistic theory that we are now devising. This is, in fact, the way Dummett thinks of operators (which is, as we have seen, the very point where classical and intuitionistic interpretation differ):

The meaning of a logical operator is given by specifying what is to count as a proof of a mathematical statement in which it is the principal operator, where it is taken as already known what counts as a proof of any of the constituent sentences (any of the instances, where the operator is a quantifier). In so far as the logical operator being explained is itself used in the explanation, the circularity is harmless, since it is a fundamental assumption that we can effectively recognize, of any mathematical construction, whether or not it is a proof of a given statement; thus, when it is explained that a construction is a proof of $\lceil A \text{ or } B \rceil$ if and only if is either a proof of A or a proof of B , the ‘or’ on the right-hand side stands between two decidable statements, and is therefore unproblematic; we are explaining the general use of ‘or’ in terms of this special use.¹²

The harmlessness Dummett is talking about is not due to the fact that the foundation of our semantics lies in some atomic statements, where logical operators would not appear. In fact, to interpret these “atoms” we still have to refer to their effective proofs, and *there* we cannot do without logical operators. Moreover, the appeal to a construction shows that, semantically speaking,

¹¹*Ibid.*, pgs. 105 f.

¹²*Ibid.*, pg. 109.

we are not dealing with atoms; we are not dealing with statements which have their own meaning inscribed within themselves. Such a circularity is nonetheless harmless, because it refers to human practice: in the example given, to the peculiar system of practices that manifests the understanding of the operator ‘or’ through the understanding of what counts as a proof of $\lceil A$ or $B \rceil$.¹³

Once again, molecularism is torn between atomism and holism, and it seems that it can escape Scylla only by giving itself to Charybdis. In fact, when we generalize the intuitionistic picture to the non-mathematical case, we see that appeals to the observation of human practice and behaviour is inescapable. We will see in the next chapter that such an appeal constitutes a founding moment for semantic holism.

It is at this point that Dummett’s hint to a revision of the linguistic practice is rendered clear. It is true, in fact, that he rightly acknowledges that his theory must take the systematic character of language into account.¹⁴ But it is also true that his proposal to view our statements as characterized by our practices in the world as more peripheric or more central with respect to the system¹⁵ (following a suggestion by Quine)¹⁶ brings some kind of metric in the system, which is, in this respect, less holistic than what we would expect, since statements are somehow weighed according to their distance from the core. On the other hand, if we take Quine’s suggestion to its end, we must acknowledge that such a metric is not unique, but that different practices at, perhaps, different times or situations may give rise to different metrics. We will come back on this subject.

About verificationist molecularism, there are two more things to be added. The first is that, usually, to know the meaning of some statement is not only to know when the statement is verified, but also to know when it is falsified, and these two things are not to be confused, because the first does not necessarily entail the other and vice versa.¹⁷

The other thing to be noted is yet another way molecularism differs from atomism. We saw how an atomistic theory which starts from the meaning of atomic sentences (as apparently molecularism does) is capable of tracing the meaning of complex sentences to the meaning of constituent (atomic) sentences. Still, there is no warranty that such an operation is legitimate. Since the meaning of a complex statement is to be traced back to an acknowledgment of what counts as verifying it, it is possible that such a verification exploits the logical operators figuring in the complex statement (which give it its complexity) in a peculiar way, which does not enable us to use that verification as

¹³ *Ibid.*

¹⁴ *Ibid.*, pg. 111.

¹⁵ *Ibid.*

¹⁶ Cf. Quine (1953), pg. 42.

¹⁷ Cf. Dummett (1976), pg. 112.

a verification of the constituent statements, and, therefore, it would not give us the meaning of the constituent statements.¹⁸

3.1.3 Meaning As Falsifying Criteria

Although promising, the solution seen in the previous section is not a definitive one. In fact, on one side, Dummett states that it is impossible, for him, to prove the truth condition solution wrong; on the other side, Dummett himself seems to prefer still a third solution.

Before introducing the third notion that might prove to be central for semantic molecularism, let us see what happened of the notion of truth. As it turns out, such notion is still somehow central, though it is no more a grounding notion for our enterprise. In fact, although we may have been persuaded that we cannot figure the meaning of statements unless we know what counts as a verification of them, nonetheless it might seem just obvious that the notion of truth is a crucial one to understand the notion of verification itself.

It is not mandatory to feel this way about truth, though. In fact, we may pursue the issue a little further, and ask whence the notion of truth itself comes from. We have already taken a route that leads us towards the analysis of behaviour – in a wide sense, i.e., of (perhaps idealized) human practice. Thus, we may feel comfortable carrying on this way, and figuring some kind of practice which would open for us a window on the notion of truth. The practice which seems more related to the notion of truth is that of making assertions.

When we think about assertions, albeit in a rather idealized way, the asymmetry between their correctness and their incorrectness strikes us. Even assuming verification as the central notion for our theory of language, we must acknowledge that recognizing a correct assertion is different from recognizing an incorrect one. In fact, recognizing that a given assertion is correct does entail the incorrectness of its negation, but the acknowledgment of the incorrectness of a given assertion does not entail the correctness of its negation.

There is, therefore, some room between the acknowledgment of the correctness or the incorrectness of a given assertion: it is the logical space where we can say that its negation is incorrect while we are still wanting for some proof of its own correctness or incorrectness. If we stick to the verificationist picture given in the previous section, we should conclude that we still do not know the meaning of the sentence expressed by such an assertion. But this does not seem to be the case:

Suppose that we are considering some assertoric sentence which we understand perfectly well in practice—that is, we have no uncertainty about the content of the assertion made by means of

¹⁸ *Ibid.*, pgs. 112 f.

it—but the application to which of the notions of truth and falsity is intuitively obscure. How do we decide whether or not any given state of affairs shows an assertion made by means of the sentence to be correct? For instance, the sentence is an indicative conditional, and the state of affairs is one in which the antecedent is recognizable false.¹⁹

According to the verificationist theory, we should conclude that, since there is no state of affairs which can count as verifying the sentence in the example, then there is no displaying of practice which can count as knowing the meaning of the sentence. This is not only counterintuitive, but also formally contradictory with the premises of the thought experiment. We are therefore forced to assume one of the following: either it is wrong to assume that we can understand a sentence even if we do not know how to verify it, or the notion of verification does not ground meaning.

To take the first option would mean to renounce too vast a part of our language – even if we limit ourselves to the assertoric use of language, as we are presently doing. The asymmetry between correctness and incorrectness of assertions, noted before, leads us into thinking that a useful notion to definitely ground meaning is at hand: the notion of falsification.

An assertion is not, normally, like an answer in a quiz programme; the speaker gets no prize for being right. It is, primarily, a guide to action on the part of the hearers (an interior judgment being a guide to action on the part of the thinker); a guide which operates by inducing in them certain expectations. [...] The expectation formed by someone who accepts an assertion is not, in the first place, characterized by his supposing that one of those recognizable states of affairs which render the assertion correct will come to obtain; for in the general case there is no bound upon the length of time which may elapse before the assertion is shown to have been correct, and then such a supposition will have, by itself, no substance. It is, rather, to be characterized by his *not* allowing for the occurrence of any state of affairs which would show the assertion to have been incorrect; a negative expectation of this kind has substance, for it can be disappointed.²⁰

I would like to underline the movement made by this progression of theories. We began with a rather transcendental notion, that of truth; then we moved on to some acknowledgment of the practical interconnectedness of sentences, and hence to the doctrine that verification (i.e., the practical acknowledgment

¹⁹*Ibid.*, pg. 123.

²⁰*Ibid.*, pg. 124, emphasis in text.

of truth) is perhaps more crucial than truth itself. Finally, again by paying more attention to the ordinary uses of language (and in particular to the rather specific assertoric use), we ended with recognizing the fundamental notion of our semantic theory in that of falsification. Thus, molecularism leads us from transcendence to concreteness, although still a rather idealized concreteness, since, after all, we are still not questioning the possibility of distinguishing assertions from other linguistic acts.

3.2 The Force of Discourses

We have seen, so far, how a semantic theory can deal with a derivation of meaning from what we could call, following the linguists, **semantic competence**, i.e., the knowledge which is implicit in a speaker's uttering a certain sentence. The connection between these two aspects of our semantic experience – the part of it concerning meaning and the part of it concerning knowledge – can be viewed under two different perspectives. On one hand, we may ask what the speaker has to know in order to express a given meaning. This was the subject of the previous section. But we may as well ask how the speaker manifests such knowledge.²¹ In other words, although our theory of meaning and our theory of interpretation are obviously intertwined, there is still a conceptual distinction that we may want to draw. It is with the latter perspective that we will be dealing in the present section.

We have seen that the determination of the meaning of any given sentence may depend on molecularistic grounds, assuming one of the three notions we analyzed as the core one. In fact, once we reduce meaning to, say, verification, the meaning of one sentence is necessarily dependent on the meaning of some other sentence stating the verifying condition of it, which we assume, for the purpose, to be known. Once we question such previous statements, though, we are of course pushed back to some other sentence, so that there is no atomic sentence, meaning by this that there is no sentence whose meaning cannot be given in terms of other sentences stating, say, its verifying conditions.

We face two problems, when dealing with such a theory. First of all, we must explain how we manage to avoid a vicious circle. Second, we must explain why we are talking of molecularism: in other words, how we can say that we do not need the whole of the linguistic system to make the semantic determination we need.

The second answer is easier to give: since we are still not dealing with an actual system, but only with an idealized one (although, admittedly, less so than that presupposed by an atomistic view), it is part of our initial assumptions that we can partition the whole linguistic system in a Quinean way. Such a partition bears great consequences on our theory of interpretation.

²¹Cf. *ibid.*, pg. 128.

The theory of interpretation may also be the key to answer the question concerning the vicious circle. In fact, if we manage to find a starting point for our semantic theory in the everyday practice of interpreting language, then, even if such a starting point is only practical and not theoretical, as long as our theory of interpretation is still molecularistic, we can see it as a support of the theory of meaning previously given. If, on the other hand, the only way to give a theory of interpretation is by some holistic means, we will be in more or less the same condition of the atomists, who cannot reconcile their theories of meaning with possible theories of interpretation exploiting the same concepts.

Michael Dummett deals with the problem of interpretation, with which we are concerned presently, when he wonders how we can derive the principles governing what counts as evidence for the truth of a given sentence from the meaning of that sentence. This sort of derivation is connected with interpretation inasmuch as the individuation of what counts as evidence for the truth of a sentence can be interpreted as a (behaviouristical or conventional) hint that what is being said is warranted by the states of affairs, by what the speaker knows, and by the semantic theory that we are exploiting to connect these to the actual semantic properties of the sentence. Moreover, once we manage to draw such a connection, we can as well follow it the other way round, going from the knowledge of the condition that allow the speaker to deem a certain sentence true to the meaning of that very sentence, thus concluding the circle. We will see that, in principle, nothing different happens when we deal with the problem of interpretation in a holistic way.

The difference, though, is in method. When we deal with the conventions governing our language and the applicability of sentences, we must acknowledge that different contexts of utterance involve different requirements for the speaker: hence, different things may count as evidence for the truth of different utterances of the same sentence.²²

This is the idea at the basis of Akeel Bilgrami's theory of meaning and interpretation. He expresses it through an externalist constraint:

When fixing an externally determined concept of an agent, one must do so by looking to indexically formulated utterances of the agent which express indexical contents containing that concept and then *picking that external determinant for the concept which is in consonance with other contents that have been fixed for the agent.*²³

Incidentally, I would like to note that, while the constraint is formulated as a principle governing some kind of theory of interpretation, it could have been stated also to be used for a theory of meaning, thus leading to a theory

²²*Ibid.*, pg. 132.

²³Bilgrami (1992), pg. 5, emphasis in text.

not much different from the Dummett's one, analyzed in the previous section. To be sure, Bilgrami adds a caveat to his constraint:

[...] although the constraint and the whole question of externalism are formulated in the context of how we "fix" an agent's concepts and contents, this does not mean that I am interested only or even primarily in the epistemological question: how do we find out about another's concepts and contents. [...] It is because I take concepts to be externally constituted that I allow myself to believe that the contents of an agent's mind are necessarily public. And it is because they are public items, items available to one another, that a scrutiny of what goes into the fixing of one's concepts and contents by another is illuminating. [...] Also, it should be obvious that when I talk of fixing concepts and attributing contents, I do not have in mind the everyday scenario where we interpret the sayings and doings of others by seeing them as contentful. I do not have in mind everyday understanding of each other but, rather, what underlies everyday understanding or what everyday understanding consists in.²⁴

In spite of Bilgrami's own purpose in devising his theory, however, I will use it as a guideline to show what a molecularistic theory of interpretation might look like. First of all, my concern will be to distinguish the externalist constraint given above from some similar constraint that may be found in some sort of atomistic theory. Even atomism, in fact, may put a stress on the role of indexicals, and then appeal to some sort of acquaintance to explain the reference of these.²⁵ But the notion of reference is not at stake here:

The centrality of indexical utterances and thoughts [...] should not give the impressions of any commitment to anything like direct reference. Indexicals are central because they are essential clues to an agent's perceptions and responses to things and events around him. This centrality is not lost if it is unaccompanied by doctrines of direct reference [...].²⁶

It is not lost because the centrality of the indexical utterances has nothing to do with direct reference. Our concern with indexicality is the same concern that Quine and Dummett had for the individuation of the (perceptual or behavioral or conventional) context of utterance.

²⁴ *Ibid.*

²⁵ Atomism has problems in explaining the interpretation of such acquaintance, of course, but the possibility of treating it as a primitive notion is not barred, though perhaps rather uninformative.

²⁶ *Ibid.*, pg. 6.

It is certainly part of the conventions governing the assertoric use of language what kind of claim we take a speaker who makes an unqualified assertion to be advancing, i.e., what kind of ground or warranty is required for the assertion not to be misleading. This is something *not* uniformly determined by the meaning of the sentence used to make the assertion, and may vary from one area of discourse to another, and also from one context to another.²⁷

The principle of externalism that we want to introduce now has to help us distinguishing the relevant characteristics of a given context of utterance, which enable us to construe the meaning of the utterance itself and, maybe, also of the sentence – i.e., of that portion of meaning that is not, in fact, dependent on the context, so that it remains the same in different contexts, if any of the like exists at all. I use the expression “context of utterance” in a rather broad sense: it is constituted by the perceptual reality which is hinted at by the indexical features of utterances, the conventional principles governing the assertability of sentences, the psycho-physical conditions which may prompt the utterance.

When we take all this into consideration, it is quite easy to see what the difference between our molecularistic theory of interpretation and a possibly atomistic theory of interpretation is. It may be more difficult, though, to see where the difference between molecularism and holism lies. We have to explain why we need not take the whole set of possible contexts into consideration to construe the meaning of a given utterance, and thus show that the differences between different contexts is so substantial that it partitions the linguistic system in different subsystems which might be considered independently from one another for the sake of interpretation.

The point, according to Bilgrami’s externalism, is that the interpretation of what a speaker means has a sense only when it is directed to the interpretation of the motives behind action. In other words, the attribution of content is propedeutic for the interpretation of human actions and, therefore, for the creation of a theory of human action and of human rationality – or, maybe, rationality *tout court*. Thus, we may need to stick to the kind of determinations that actually prove useful with regard to this task, and these are such that we need not be concerned with the whole systematic nature of language.

Action-explanation always takes place at a much more local level than the meaning-theoretic level. Here, the entire aggregate of beliefs that an agent associates with “water” are not all relevant. One distils out of the aggregate of resources provided by the meaning-theory only those beliefs that are relevant to the action-explanation at the local level. Thus, if you and I are both drinking some substance from the kitchen tap because we want to quench our

²⁷Dummett (1976), pg. 132, emphasis in text.

thirst with the cheapest available drink, we may *in this locality* both be attributed the same content: "...that water will quench thirst." In this locality of explanation, your chemical beliefs are simply not among the beliefs selected from the specifications of the meaning-theory (for your idiolect). The specification of your local water-concept in the attribution of content which explains your behavior does not contain your chemical beliefs. Those beliefs are not needed in this locality so the specification of the local water-concept ignores them.²⁸

The major objection that holists may oppose to such a theory is that it presupposes the whole semantic system as a pool from which local contents are distilled. Bilgrami grants it, but he thinks that it is no victory for a traditional holistic view, since that role is almost influential.²⁹ Every pool of contents, in fact, can do the work, and any explanation for the process of arriving to a local set of concepts is good for our purposes.

If this is correct, we may as well think that the notion of a whole pool of meanings from which we can draw the information we need with the help of contextual information is a derivative one, and not a primitive one. In other words, we might start with the interpretation of individual discourses – i.e., groups of utterances connected to the same locality. Such a proposal calls for an explanation of the recognition of the same content in different localities. Of course, our answer should be that it is not appropriate to speak of exactly the same content, but we may be nonetheless be interested in determining if we are dealing with connected meanings of the same word or utterance, or if we simply stumbled in an ambiguity of the language. Given the two localities of Bilgrami's example (drinking water from the tap, and talking of water in a chemical lab), we want to know if the water-concepts which they individuate are different like the bank of the river is different from the money bank, or rather like a carnival mask is different from a gas mask.

Bilgrami would like to redimension this problem as well. He does it by looking at a connected question, the dichotomy between change of meaning and change of theory. Another way to look to the problem we posed in the last paragraph, in fact, is to look at it from the perspective of a change of opinion about what constitutes the concept of "water." Such a change can be motivated by very different causes, of course: mere scientific progress or correction of a linguistical imprecision being only two extreme instances. We want to know, in the new formulation of the problem, whether ancient Greek scientists had the same concept of water that we have, and only a different scientific theory built around the concept (hence, a difference in beliefs), or our concept and theirs are altogether different.

²⁸Bilgrami (1992), pg. 11, emphasis in text.

²⁹Cf. *ibid.*, pgs. 151 f.

If my picture of things is correct, then meaning is specified by those beliefs that are invariant between changes (or disagreement) in belief. And that I find a very intuitive picture of what one notion of meaning is. This notion of meaning fixes the subject for belief-change. If there is nothing invariant then there is no preservation of meaning, and the change must be counted as a case of meaning-change.³⁰

The question, then, is again about the construction and the individuation of the relevant localities and, therefore, of the relevant overlapping beliefs. Bilgrami thinks that, since the process of interpretation is necessarily open to revision, we only need a rough starting point, which can as well be furnished by some intuitive theory.³¹

3.3 Evaluating Semantic Molecularism

As in the previous chapter, I will devote a section to the examination of flaws or obscure points in the semantic theory sketched in the previous two sections. Such flaws will be divided in two categories: those that arise by merely considering the theory as it stands, and those that have their origin in the interaction between semantic molecularism and other theories or disciplines.

3.3.1 Internal Problems

We have already examined the reasons that led Dummett shift from a theory of meaning grounded on the notion of truth to a theory grounded on the notion of verification, to one grounded on the notion of falsification. The faults of the two first kinds of semantics, then, should by now be obvious enough.

In a nutshell: we may find problems in a truth-based theory of meaning because the notion of truth requires some knowledge of it; hence, either we consider truth and verification as interchangeable notions for our purposes,³² or we assume some sort of faculty (acquaintance) to connect us with truth, which must be different from the verification faculty. As for verification, we can see from rather common examples such as indicative conditionals with a false antecedent, that we may very well understand the meaning of an assertion, even without knowing what may count as a verification of such assertion.

³⁰*Ibid.*, pg. 128. Cf. *ibid.*, pgs. 152 f.

³¹*Ibid.*, pg. 146.

³²Notice that this is not equivalent to treat the two notions as interchangeable in all cases, because some non-linguistical theory may help us to distinguish them. Still, if we want to save some kind of founding role for the notion of truth, we will have to prove that, in the context of a given theory or discipline, truth is not only not interchangeable with, but also not derivative of verification.

Before dealing with falsificationism, let us face a quite obvious spot which flaws any theory starting with some kind of access to truth conditions, viz., the fact that the theory necessarily deals first of all with statements and assertions. I would like to make it explicit that, by doing so, a whole section of the theory is still missing, the part explaining how to construe the meaning of sentences with a force other than the assertoric. That may not be a difficult task, and we indeed have some hint of how to proceed. Some care needs still to be taken, because the fact that in some case we find it quite easy to reconnect to the assertoric case does not entail that this is as easy or even as possible for any kind of sentence or utterance.

We made a similar point about atomism,³³ but here the difficulty is not as fatal, because we are allowed to consider sets of utterances or sentences as aggregates around a practice, and this can show us a way for such an enterprise. We should consider this as an occasion for further research, rather than an unsurmountable flaw.

When dealing with the ultimate theory of meaning devised by molecularists, we must remember that such a theory is not distinguishable, for what concerns its merits, from the theory of interpretation that accompanies it. The major objections against this is that it is still not clear whether we have the right to talk about molecularist interpretation at all. In Bilgrami's theory (which he does not call a molecularism) there is an important role played by the pool of the whole aggregates of meaning that is undeniable. In spite of Bilgrami's attempts to belittle its contribution, it is very difficult to make the theory work without considering the holistic meaning system original. But if we do such a thing, then we are dealing with a holistic theory of interpretation, and all the talk about localities merely describes the practical (though still somewhat idealized) application of the general theory. I do not want, on my turn, to belittle such a theoretical contribution, because it is of undoubtful importance to show how a theory *prima facie* implausible – as holism is – might work in practice. Nonetheless, if we devise a theory to handle holism and the infinity of possibilities that holism gives us, we still end up with holism and not with a molecularistic theory of fractionary infiniteness.³⁴

This we would have, if we could show localities to be original, and the holistic semantic system to be derivative. I do not think that Bilgrami had anything the like in mind, and nothing in his book suggests it anyway. But the proposal is clear enough to be taken into consideration.

Unfortunately, it is not as clear as to be spotless. If we have to consider localities as original, it is legitimate to ask how the beliefs which are relevant for

³³Cf. § 2.3.1 above.

³⁴More on the subject of holism, the theoretical possibilities it furnishes, and ways to handle them, will be said in the following chapters.

interpretation in localities are individuated. Bilgrami dismissed the question as follows:

The protest that there is no principle for deciding what the relevant beliefs are in the local level is [...] exaggerated. All one needs to do in order to fasten on the right explanandum and the right explanans, at the local level, is to see what beliefs the two agents will agree on. Imagine them communicating in this local context of drinking water and see what they agree on and what beliefs puzzle and throw them off. If one of them says that the substance that they wish to drink will not poison them and the other agrees, then that belief may be counted as relevant to this local concept “water” which goes into the explanation of their actions. If one of them says that the substance has the chemical composition H_2O and the other is puzzled by that, count that belief as irrelevant in this locality. There is no serious obstacle to finding out which the relevant and irrelevant beliefs are.³⁵

This appeal to the intersection of belief seems to me more problematic than Bilgrami takes it to be. It is clear that for the example to work, we can take into consideration all the possible pairs of thirsty people who meet at a faucet. If our test takes into consideration two chemists, then beliefs about chemical composition of drinking water become relevant. This might imply that we have a different locality every time a situation occurs: in other words, there are no two equivalent situations. This begs the question regarding the individuation of localities. I will deal with this in a moment, but I would like to point to another problem about the intersection theory before.

Consider what our situation is: we need the theory of localities because we want to attribute content to people in order to explain their action, and people act in localities. Now, according to the intersection theory, we have the problematic view that a chemist drinks tap water for different reasons depending on the presence near the faucet of another chemist rather than a person ignorant of chemistry. This is not only counterintuitive, but also inconsistent with our claim that such a theory would help us explaining human action: it seems, if anything, that things get more confuse.

Maybe we could interpret the intersection theory in some other way. For instance, we could imagine that drinking tap water is a locality that we can individuate, and that the relevant beliefs we are looking for are given by the intersection of the beliefs of any two people who could participate in that activity. Or, which would amount to the same result, we could imagine that all the people met at the same faucet to quench their thirst. Then, the beliefs which are relevant to explain their action will be those that any single one of them has. But we have a problem with this generalized intersection theory:

³⁵Bilgrami (1992), pg. 146.

As a matter of fact, each of us knows or believes partly different things (partly different subset of the collective encyclopedia), yet we are all competent in the use of our language. Could lexical competence be defined as the *intersection* of such different individual competence? In principle, no, for we have no guarantee that the intersection of all competences is rich enough to be considered adequate as a competence (it is quite possible that the set of beliefs we all share about, say, gold is very small, too small to constitute an adequate competence with ‘gold’).³⁶

A defender of the intersection theory may object that here Marconi is dealing with something completely different, because the competence with a lexical item is not at stake when we are trying to explain the behavior of people acting in a certain way. For, if people did not share at least some relevant and important belief about tap water, they would not drink it to quench their thirst; hence, we would have no behavior to explain.

The objection is fundamentally correct, but for three gaps that need filling. One is again the question about the individuation of localities, and I still propose to leave it aside. The second one is that we are completely begging the question against the existence of a core water-concept in the tap locality, if we reason that, since people drink water from the tap to quench their thirsts, then there must be a unique reason (and a unique concept underlying that reason) for everyone. If these are our only argument, we still cannot bar that any single human being could drink water from the tap for entirely different reasons, and with entirely different concepts in mind.

The last gap is connected to the problem of lexical competence. We might agree that the intersection theory is to be used only with some locality in the background, but we nonetheless might ask the meaning of a word in general, cross-locally. Here, intersection is not working any more. So we are driven to the conclusion that the meaning attached to some lexical item may vary from individual to individual, and even for the same individual across time.³⁷ There is nothing wrong with such a conclusion, as long as it is acknowledged that it must be grounded on a theory of localities: in order to accept the variance of meaning depending on the situation we need to know how to distinguish when situations themselves vary.

We have seen that a possible answer is to start by appealing to our intuition of what constitutes a locality, and then revise our theory in order to make sense of inconsistencies. The problem with such an answer is that if we want to consider the discourse about localities as original we must explain whence our intuitions about them come from.

Let us assume that such knowledge may come in a molecularistic way. What we need, then, is some more original locality (or molecule) to guide us

³⁶Marconi (1997), pg. 52, emphasis in text.

³⁷Cf. *ibid.*, pg. 54.

toward the choice of the locality which is needed to interpret a certain piece of human behavior, e.g., an utterance. I think this can be done only by taking into consideration the context of uttering, and our beliefs about that uttering and its context, and so on. To choose the right molecule to interpret our original utterance, thus, we end up with a bigger molecule. To know if this bigger molecule leads to the right interpretation, however, we may have to compare it with some other molecule. Eventually, we might need to consider the whole language.

In this sense, I think that a molecularistic theory of interpretation is not different from a holistic theory of interpretation, and there is a reason for that: molecularism, to be implemented, needs the possibility of revision that only a holistic picture entails. Of course, we might think that we are dealing here with some different kind of holism than the traditional one. This is, so to say, a potential holism, as opposed to an actual one. In the picture given here, that is, we merely consider the possibility of drawing connections between the various elements of the system, while more traditional holism considers such connections as already given. This may be true: however, all we wanted to show was that we cannot consider the whole totality of the system as derivative, but that we need to consider it as being on the background of every decision we make in the process of interpretation.

3.3.2 Interactions

The disciplines that look more related to the issues dealt with by semantic molecularism seem to be linguistics, psychology and epistemology. Contrary to semantic atomism, molecularism does not have great ontological or metaphysical problems. The shift from the notion of truth (which is a metaphysical notion) to the notion of access to truth – either by means of verifying conditions or by means of falsifying conditions – closes the door in the face of the metaphysical debate, only to open it to the epistemological one.

While verificationism has its flaws, as we have seen, it seems that falsificationism has a better fate. It must be clear that, in spite of its apparent affinity with Popper's philosophy of science (which is, as it is known, founded on the notion of falsification), Dummett's theory does not seem to share the faults of Popper's. There is, perhaps, some problem in analyzing in terms of falsification probability statements, though – which is indeed one of the objections that have been risen against Popper.

Also, we must note that the notion of falsification involved here is somewhat laxer than Popper's, for the philosopher of science wanted it to function as a demarcation criterion between science and non-science, while we want it to be an effective means for construing the meaning of any kind of assertion, whether scientific ones or not. I will not dilate on the subject of demarcation: suffice it to say that it is not at all clear whether such a thing as a demarcation criterion between scientific and non-scientific statements exists at all, let alone

whether falsification can be such a criterion.

It should be clear, on the other hand, that applying a falsification theory of meaning to, say, metaphysical statements involves a rather complex theory of expectations. We would have to state, in fact, what is the warranty that we may have to have certain expectation with respect to things or events that we know are beyond our capacity of surveying, and in what way such expectations affects what we are legitimate to assert concerning those things or events. In this sense, it is obvious that epistemology should deal not only with science (even in a broad sense of the term), but with all the many forms which human knowledge may assume.

This task, therefore, borders psychology at least in the minimal sense that is human knowledge and human activities that are involved. After all, as we have said, molecularism offers itself as a theory primarily aimed at the explanation of human action. Some methodological question about psychology, thus, concerns also molecularism. The most important of these is perhaps a question about justification: whether we are justified in assuming the existence of certain invisible and inscrutable processes by the fact that a theory exploiting them entails certain phenomenic consequences which we indeed make experience of. There is a large debate about this issue (which does not limit itself to psychology or human science, but involves even hard sciences such as physics), and several kinds of answer might be given.

Psychology, in its turn, is connected with neurosciences, and an answer to our methodological question might come from results of the latter discipline, depending on our opinion on what the actual relationship between the two is or should be. This is a question concerning theoretical reduction, and again it is not an unusual one even outside the domain of cognitive science.

Then, there is the question concerning linguistics. The stress put by molecularism on human practice makes pragmatics the ideal interlocutor. After all, some pragmaticists would be happy to have a philosophical background to help them in the individuation of the relevant localities and contexts for the interpretation of utterances. Fixing the flaws of the molecularistic theory of interpretation would maybe help an advancement of pragmatics; and the opposite kind of influence cannot be excluded, either.

3.4 Some Brief Conclusion

I will not hide my sympathy for molecularism. The stress on human activity, the intrinsic dynamism that is attributed to the process of interpretation and language use, the avoidance of suspicious metaphysical notions such as that of truth are qualities that I find desirable for a theory dealing with human language. Still, these are not universal desiderata, since it is not obvious that our first concern should be with human language and not with something else, perhaps language per se.

One important difference between semantic atomism and semantic molecularism is, in my opinion, that the former succeeds in giving the answers to the kind of questions it is interested in, while the latter's flaws may seem to prevent it to do it. Moreover, semantic holism will be shown to deal with the same kind of questions; therefore it will be a good thing, for us, if the desiderata fulfilled by molecularism are not forgotten by holism. In this sense, besides being a competitor theory (perhaps even more so than atomism), molecularism should be deemed a great ally of holism.

Part II

Semantic Holism

CHAPTER 4

Holistic Interpretation

Now we turn to the topic of semantic holism. In the previous chapters we have said that the fundamental problems of philosophy of language are: 1) explaining how we understand each other when we speak, and 2) how – or in what sense – our language can be about our world. We have also seen how the atomistic strategy of answering first to the second problem fails to indicate a way to answer the first. Semantic holism attributes this failure to the wrong order in which atomism answers the questions: one should reflect first on communication between human beings, and only after that on the connection between language and world. This is the strategy I will follow.

We have seen in § 1.2 that the characterization of languages in terms of sets of rules – which is the characterization we have chosen to adopt – has as a consequence the attribution to each different speaker of a different language. The problem of interpreting others' speech, thus, cannot in principle be different from the problem of translating from one language to another. Our concern, thus, is to find a good theory to explain why given sentences in one language correspond to certain sentences in another and not to other ones.

Semantic atomism made an assumption about interpretation we may find rather natural and commonsensical. The distinction between semantic production and semantic interpretation implies that all interpretation is indeed interpretation of meaning *meant by someone*, who must, therefore, produce it first. The assumption was, thence, that *semantic production comes before semantic interpretation*. Therefore, semantic atomism deems semantic productivity (i.e., the ability of producing meaning) the essential feature of language. Hence, the relative comfortableness with which it gives up to a partly holistic strategy to explain interpretation: the latter is not the main aim of an atomistic theory, and can be left to other theories, perhaps even to other disciplines altogether, such as psychology.

Holistic semanticists hold a different view about language. According to them, it may well be true that the most important feature of language is its

semantic productivity. Still, this is certainly one of the most difficult features to explain, and one whose mechanisms seem most hidden. Therefore, it might seem more profitable to begin our investigation of the semantic features of language from a phenomenon that, although it may seem a minor feature at first sight, is nevertheless evident in every single use of language, and not less mysterious and requiring philosophical investigation than semantic production. The feature I have in mind is of course interpretation.

Production, thus, will be addressed by holistic semantics only after having coped with interpretation. The holistic mechanisms uncovered by an analysis of this linguistic task as an original task – as opposed to a secondary one – will furnish us with new means to find a unitary solution to both the fundamental questions about language we are asking ourselves. Before analyzing in depth the issue of interpretation, though, it is necessary to clarify the notion of holism itself. This will be my starting point.

4.1 The Notion of Holism

The notion of semantic holism, as it may appear obvious, depends on what we understand by “meaning.” It is important, for this discussion, that the difference brought in § 1.3 is a clear one. It is different, and it makes a difference for our discussion, if what we are concerned with is study of reference or of meaning in the proper sense. And in both cases we have to distinguish between semantic properties as properties of a system (or of parts or elements of a system, since this distinction comes before any argument in favour or against semantic holism) or as properties of people. It is different, in other words, to be concerned with those features of language that make it possible for words and sentences to have the semantic properties that they have, on one side, and, on the other side, to be concerned with the way those properties make it possible for us to communicate and interact with other people.

4.1.1 About an Eccentric Definition of Holism

A theory can be holistic with regard to all of these different aspects of “meaning,” and we must bear in mind this when evaluating holistic theories, because we might risk to impute flaws to an argument only because we misunderstood the actual aim of the theory under examination. Fodor and Lepore are amongst the strictest critics of semantic holism: they believe that no good theory has been put forward to prove semantic holism.¹ In their book, they propose two different notions of holism. One can be found in their preface:

This is a book about holism about meaning: roughly, it’s about the doctrine that only whole languages or whole theories or whole belief systems *really* have meanings, so that the meanings of smaller

¹Fodor and Lepore (1992), pg. xii.

units – words, sentences, hypotheses, predictions, discourses, dialogues, texts, thoughts, and the like – are merely derivative.²

This definition makes explicit reference to systematic wholes as being the units where meaning is determined. What is essential in this definition is not only that we consider meaning as a property that originally pertains to wholes, but that these wholes are considered as structured. This can be an essential difference, because we are accustomed to think of languages as infinite structures and our definition must therefore account for these. If the elements we are interested in are sentences, hypotheses, predictions, discourses, dialogues, texts, thoughts and the like, our interest resides primarily in systems which allow an infinity of different such elements.

I am not claiming that, according to this view, only infinite systems can be regarded as languages: what I am claiming is that only with respect to infinite systems there is a point in talking of languages.³ If this is correct, we must be interested in some regularity to be found in the systems we want to study: otherwise, there is no chance that we can end up with a theory about infinity. Thus, saying that holism about meaning is the doctrine that only such structured systems have meanings (as Fodor and Lepore claim) is equivalent to saying that it is the doctrine that meanings are produced only through certain kinds of systematic rules.

Fodor and Lepore proceed to distinguish holism about meaning from several connected issues:

Meaning (or semantic) holism must be distinguished from a number of related ideas with which is easily confused: from holism about confirmation, for example, or about interpretation or about the individuation of functional properties. These other kinds of holism might be true even if meaning holism isn't. Or at least, so it seems; unless there are arguments to show not just that confirmation, interpretation, or the individuation of functional properties is holistic but also that the meaning of a symbol is somehow

²*Ibid.*, pg. x, emphasis in text.

³A counterexample might be seen in Wittgenstein's linguistic games, and toy-languages, such as can be found in Wittgenstein (1953), § 2 (cf. Fodor and Lepore (1992), pg. 234 n. 12, where the Wittgensteinian "primitive language" is used as a counterexample to a different, but related, point). I think that, even if Wittgenstein thought that such languages could be imagined as "the *whole* language [...] of a tribe" (Wittgenstein (1953), § 6, emphasis in text), they actually deserve the name of languages only because of a comparison with more proper languages such as our own. Moreover, consider that Wittgenstein's thought experiment is, of course, set up starting with a portion of our language; maybe, then, our impression that it is still a language is biased from the beginning because of the nature of the thought experiment. Finally, the decision that we may make to consider those speakers as members of the human species can play a role in our attributing them a language and, therefore, in our considering the Wittgensteinian toy as a language. This last point will become clearer later (see §§ 4.2.7 and 4.3.2).

determined by facts about how its applications are confirmed or how its tokens are interpreted or the functions that it performs.⁴

There is a sense in which all that counts about symbols is indeed interpretation, confirmation and individuation of function. It is in this sense that Ludwig Wittgenstein claimed that meaning is use.⁵ According to this view, there is nothing more in the meaning of a symbol than the way the symbol is used. Since this use is in turn regimented by some system of rules that regulates all uses of all symbols in the symbolic system, each in relation to the others, claiming that meaning is equivalent to use is a way of holding a holistic theory about meaning.

Fodor and Lepore have problems with this view being considered as genuine semantic holism. To see what the problem is, though, we have to introduce the second definition they give of the term ‘holism,’ which is the one that they assume as correct and discuss in the rest of the book. This definition, in turn, comes together with definition of other two terms, viz., ‘atomistic’ and the rather infelicitous ‘anatomic’ (which should be read in the meaning of ‘not atomic’):

Many properties have the property of being, as we shall say, *anatomic*. A property is anatomic just in case if anything has it, then at least one other thing does. [...]

If a property is not anatomic, then we shall say that it is *atomistic* or *punctate*. An atomistic or punctate property is one which might, in principle, be instantiated by only one thing. [...] One way of formulating a main issue to be discussed in this book is whether *being a symbol*, *being a symbol belonging to language L*, *having an intentional object*, *having intentional content*, *expressing a proposition*, *having a referent*, *being semantically evaluable*, and the like are punctate properties. The currently received philosophical view is that these sorts of properties are not punctate but anatomic. [...]

Many anatomic properties have the property of being *very* anatomic, or, as we shall say, the property of being *holistic*. Holistic properties are properties such that, if anything has them, then *lots* of other things must have them too. The “lots of” part of this definition could bear to be sharpened, no doubt; but, for our purposes, this isn’t required. Our primary concerns in this book will be with natural languages and with minds. Natural languages and minds can be assumed to be productive in all the interesting cases; [...]. The semantic properties we’ll discuss will therefore

⁴Fodor and Lepore (1992), pg. x.

⁵Wittgenstein (1953), § 41.

generally be ones which, if they are holistic, then if anything at all has them, so too do endlessly many other things.⁶

From this general definition of holistic properties, Fodor and Lepore distinguish two related issues about semantic holism:

What we will call *content holism* is the claim that properties like *having content* are holistic in the sense that no expression in a language can have them unless many other (nonsynonymous) expressions in that language have them too. In effect, it's the doctrine that there can be no punctate languages. What we will call *translation holism* is the claim that properties like *meaning the same as some formula or another of L* are holistic in the sense that nothing can translate a formula of L unless it belongs to a language containing many (nonsynonymous) formulas that translate formulas of L.⁷

Let us consider the difference between this definition and that given before. If we think that semantic holism has something to do with the number of entities having semantic properties, then it is obvious that finite languages or punctate languages, if possible at all, cannot be holistic. On the other hand, if what matters is only the systematicity of the whole structured set of entities under inquiry, then the number of them is not in question, and finite and punctate languages (again, if independently proven possible) can count as holistic languages. A punctate language, in this respect, will be a degenerate case, in which the entity under inquiry actually is the whole system, and therefore it is trivially true that its meaning depends on (actually is coincidental with) the meaning of the whole system.

Fodor and Lepore's mistake seems to be that they consider what we are interested in (natural languages and human minds) as the only possible subject of investigation. Thus, they cut a definition which is good enough to characterize these subject,⁸ but which could need sharpening and rethinking, if meant to be a general definition. On the other hand, the former and more traditional definition was good enough to characterize every system we could think of as endowed of semantic properties. Of course, we are interested mostly in the study of infinite systems, such as natural languages and human minds, but we must decide whether is in our intentions to consider also other systems or not. In case we decide that we are interested in other systems as well, our definitions must allow the consideration of such systems since the beginning, so that we can use them as counterexamples only if the notion of

⁶Fodor and Lepore (1992), pgs. 1 f., emphasis in text.

⁷*Ibid.*, pg. 5 f.

⁸Indeed, for such cases, it can be proved that Fodor and Lepore's definition of holism is equivalent to the traditional one; cf. Marconi (1997), pg. 49.

holism, while in principle applicable, is actually proved not to apply. If, on the other hand, we are not interested in finite systems, our definition can be cut to leave them out, and thus we cannot take such systems as counterexamples.⁹

Why do Fodor and Lepore give their definition instead of a more traditional one? Two reasons can be found in what they say. The first is that the notion of anatomisity, in their opinion, proves useful also in other philosophical contexts. The example that they give is the argument against the possibility of a private language.¹⁰ I do not want to enter into details here (and even Fodor and Lepore consider the question as a rather marginal one), but I would emphasize that to me is not at all clear how mentioning anatomisity helps “to distinguish the part of the private language problem that’s about language from the part that’s about privacy.”¹¹ I feel that it was clear even without introducing fancy anatomisity that the argument was about “the conceptual possibility of a language with only one speaker.”¹²

The second reason is more compelling: Fodor and Lepore claim that there is an argument from the anatomisity of semantic properties to semantic holism. They add that versions of such an argument are the ground of all discussions of holism.¹³ Thus, it seems that the notion of anatomisity was already beneath most arguments about semantic holism, and that Fodor and Lepore only uncovered it.¹⁴ Let us then examine this argument, and what Fodor and Lepore have to say about it.

4.1.2 The Argument A

Fodor and Lepore hold the classic argument for semantic holism – which they call “Argument A”¹⁵ – to go along these lines:

- Premise 1:** Generic semantic properties like [...] being-some-or-other-belief-of-Smith’s, being a formula of language L, etc. are anatomic. [...]
- Premise 2:** There is no *principled* distinction between the propositions that Smith has to believe to believe that P and the propositions that Smith doesn’t have to believe to believe that P. [...]
- Conclusion:** The property of being-some-or-other-belief-of-Smith’s is holistic.¹⁶

⁹For a general perspective on the issue of finite languages, see § 1.2.

¹⁰Fodor and Lepore (1992), pg. 6.

¹¹*Ibid.*

¹²*Ibid.*

¹³*Ibid.*, pgs. 22 ff.

¹⁴Cf. *ibid.*, pg. 26 for a list of arguments on the lines of the one they reconstruct.

¹⁵*Ibid.*, pg. 23.

¹⁶*Ibid.*, pgs. 23 ff., emphasis in text.

Let us note that, as Fodor and Lepore stress, the reference to Smith is inessential and, also, that the argument could hold even if we substitute different intentional properties to beliefs.¹⁷ In fact, as it stands, the argument is supposed to prove the impossibility of a punctate mind, i.e., “a mind which can entertain only one proposition.”¹⁸ But we can change Premise 2 so that it reads: “There is no principled distinction between the meaningful sentences that language \mathcal{L} must allow in order for a given sentence s to be a meaningful sentence of \mathcal{L} ” (call this Premise 2*). Through Premise 2* we can reach Conclusion*: “The property of being a meaningful sentence of \mathcal{L} is holistic,” which is the same as claiming the impossibility of a punctate language. With opportune changes to Premise 2 we can reach analogous conclusions for all the semantic or intentional properties we might be interested in.

There are three orders of consideration to be done about Argument A. We can question the validity of one or the other of the premises, or of the inference from the premises to the conclusion. Fodor and Lepore’s strategy throughout their book is showing that there is no argument for Premise 1.¹⁹

The status of Premise 1, in its turn, depends strictly on the intentional property we are concerned with. Therefore, an evaluation of the premise can be given only by way of a thorough holistic theory of the intentional property of the case. So, we should postpone our analysis of Premise 1 until further sections, and concentrate, for the present, on Premise 2 (or 2*, etc.) and on the status of the inference.

The Analytic/Synthetic Distinction

Premise 2, in all the forms it can take, is a form of rejection of a distinction between analytic and synthetic sentences (a/s distinction).²⁰ The reading of the rejection of the a/s distinction involved in the premise is “The a/s distinction isn’t principled.”²¹ The first source of an argument in favour of the rejection is, of course, Quine’s “Two Dogmas of Empiricism.”²² However, Fodor and Lepore claim that Quine’s argument is incompatible with Argument A:

When Quine says “No a/s,” he presumably means “No analytic sentences.” On that reading of premise 2, however, argument A would appear to be inconsistent. For premise 1 requires that there be sentences other than P that must be believed if P is believed;

¹⁷ *Ibid.*, pg. 25.

¹⁸ *Ibid.*

¹⁹ *Ibid.*, pg. 31.

²⁰ *Ibid.*, pg. 24.

²¹ *Ibid.*, pg. 25.

²² Quine (1953), chpt. 2.

and it looks as though the hypothetical formed by writing one of these sentences after “if P, then” must be analytic [...].²³

There are two claims that are being made in this quotation: one is a matter of Quinean exegesis, the other is about the correct reading of Premise 1. Let us deal with the former, first.

It is definitely true that Quine was more concerned with analyticity than with syntheticity in his paper. After all, his problem was with this notion and with the idea that there are truths never to be revised that can be well-distinguished from alleged truths that may prove false in face of some evidence. In spite of this, I cannot see the distinction that Fodor and Lepore would like to draw between Quine’s position and the one entailed by Argument A. If you believe that there are no analytic sentences, you also believe that the a/s distinction is not principled and that, as a matter of fact, analyticity is not a property that sentences have on their own, but only with relation to the place they occupy inside a theory. In other words, since theories are developed by human beings for their purposes, it ultimately depends on human purposes and on human decision which sentences are to be considered analytic at a given time:

Furthermore it becomes folly to seek a boundary between synthetic statements, which hold contingently on experience, and analytic statements, which hold come what may. Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system. Even a statement very close to the periphery can be held true in the face of recalcitrant experience by pleading hallucination or by amending certain statements of the kind called logical laws. Conversely, by the same token, no statement is immune to revision. Revision even of the logical law of the excluded middle has been proposed as a means of simplifying quantum mechanics; and what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle?²⁴

Quine’s point does not seem to be simply that there are not analytic sentences, but rather that their individuation is not principled, since it happens for pragmatical reasons. It could be objected that the definition of analytic statements given in the above quotation is not the standard one. Rather, it is a surrogate accepted by Quine after having refuted all more standard definitions. In fact, it is not so: Quine was quite explicit in stating that giving a definition of analyticity in terms of condition of verification (viz., a statement

²³Fodor and Lepore (1992), pg. 25.

²⁴Quine (1953), pg. 43.

is analytic when it is true come what may) is equivalent to standard definitions of analyticity in terms of identity of meaning.²⁵ Thus, I am claiming that Quine did not show that there are no analytic sentences at all, but that we cannot define analyticity, unless we give a definition which is bounded to make the individuation of analytic statements and sentences dependent on our pragmatical reasons. In other words, there is no principled a/s distinction.²⁶

Note, furthermore, that even if we agreed with Fodor and Lepore's account of Quine's position, that would leave open the possibility of building Argument A. Premise 1, which is allegedly incompatible with Fodor and Lepore's reading of Quine, is in fact perfectly consistent with it. For let us assume that there are no analytic sentences. What Premise 1 is saying is that we can build Fodor and Lepore's implication "if P , then" and hold it as immune to revision for our practical reasons. This is, allegedly, Quine's surrogate for analyticity, and that can be used in exactly the same way in Premise 1. What premise 1 says is, e.g., that in order to believe P there must be some Q that you also believe. Premise 1 does not say that Q must be determined, and that it cannot change across time. The same holds, *mutatis mutandis*, if we build Premise 1 around a semantic property like "being a meaningful sentence of language \mathcal{L} ." Thus, an evaluation of Argument A as a complex would bring more to the interpretation of Premise 1.

The Inference from Anatomisity to Holism

Let us now deal with Argument A's inference form, then. Fodor and Lepore rightly note that this argument is a version of the sorites paradox.²⁷ This, however, should not automatically imply that the conclusion is false. All that is implied is that the distinction here at stake – between, say, which meaningful sentences \mathcal{L} has to allow in order for a given sentence s to be a meaningful sentence of \mathcal{L} , and which it does not have to – is arbitrarily traced, if traced at all. Before such distinction we cannot know the way sentences in a language are connected. Moreover, even when such distinction is in fact drawn, we know that we can always revise it. When we change the criteria to trace the distinction between sentences related with a given one and sentences unrelated with it, also the boundary changes. Nothing excludes that there could be criteria which connect all sentences to the given one.

It might be objected that it does not matter whether we are capable or not of drawing the distinction, for it may nonetheless exist. In other words, the argument could be thought of as showing that we must be holists when collecting our semantic data, but that semantic properties and semantic objects can be arranged in a non-holistic way. I respond that we should always be

²⁵ *Ibid.*, pg. 38.

²⁶ More on the subject of analyticity in §§ 4.2.2 and 5.2.4.

²⁷ Fodor and Lepore (1992), pg. 25.

suspicious about an argument that says that, no matter what our theories are, there is always a chance that the real world is different from how we picture it. If this is supposed to show that all knowledge is revisable, I firmly agree. If it is supposed to prove the existence of a real way in which things work as opposed to the way we figured them to work, let me ask on my turn what this real way has to do with us. If we claim that all our philosophical or scientific theories will always be incapable of tracing such a distinction, why should we believe in the real existence of such a distinction?

I want to stress that although the form of the argument is a sorites, it is not the classical sorites by accumulation (like the classical example of baldness and hairiness, for instance). Our problem resembles rather the problem of identifying the border of a mountain. It resembles it in the sense that probably there is no way to draw the border of the mountain so that it respects the real border that the mountain actually has. In fact, there are reasons to think that there is no such thing like a real border.

There is a difference between the two cases, though, because in the case of the mountain we know we want to draw the border somewhere, because we know that the territory for which it is meaningful to ask whether it is part of our mountain or not does not extend infinitely. We also know that we do not want all the territory under consideration to end up as being part of the mountain. In the case of our semantic system, we are quite sure that it does extend infinitely, and we have no principled bias against its entirely belonging to one side of the border. Indeed, if we had such a bias, the whole discussion concerning holism would be pointless. Thus, Argument A amounts to saying that, (Premise 1) since every part of the infinite territory surrounding (so to say) a particular semantic element e (a meaningful sentence, or a belief, etc.) might prove relevant to the determination of e , and (Premise 2) we have reasons to believe that we cannot determine any subterritory that is the only relevant in all circumstances, then (Conclusion) all the system is in fact relevant to the determination of e .

Fodor and Lepore try to give strength to their point against Argument A with another argument:

Imagine that there are disjoint sets of propositions such that (1) believing any of these sets is sufficient to be able to believe P; (2) you must believe *at least* one of these sets in order to believe P; (3) none of these sets is such that you must believe *it* in order to believe P. So someone can believe P if he believes A, *or* if he believes B, and so on. [...] Then, on the one hand, premise 1 would clearly be true. And premise 2 would be true in at least the sense that there are no analytic beliefs. Yet neither content holism nor translation holism would follow. Content holism wouldn't follow because it requires that there must be *many* other propositions that I believe

if I believe that P [...], and the current assumptions allow that some or all of the disjoint sets each of which is sufficient to be able to believe P might be quite small. So, compatible with the present account, content might be molecular rather than holistic.²⁸

This argument seems to me to have a flaw. Let us assume that A (or B , or whatever) is actually the set of beliefs that allows Smith to believe P . Now we can ask why Smith believes A . Remember that Fodor and Lepore claim that A can be a very small set. They also claim that the argument allows a reading of the rejection of the a/s distinction according to which there are no analytic sentences.²⁹ So we can ask, about any of the few beliefs in A what is sufficient, or necessary, to believe them. The question is legitimate, unless we claim that the beliefs in A are self-sufficient. And they cannot be self-sufficient, unless we contradict Premise 1. So, for every belief in A , we are in the same situation as with P .

Now, P cannot be sufficient to believe any of the beliefs in A , because otherwise A would result necessary to believe P , against hypotheses. Also, given a belief a in A , the other beliefs in A cannot be sufficient to believe it, otherwise they would also have been sufficient to believe P in the first place. Suppose, then, that for each belief a_i in A we can individuate the set B_i which is responsible, being sufficient but not necessary, for Smith believing it. For every belief in every B_i we are again in the same situation as P and the beliefs in A : again P cannot be sufficient to believe any of the beliefs in B_i , because the transitivity of implication gives $P \rightarrow b_{i_j} \rightarrow a_i \Rightarrow P \rightarrow a_i$, and we already excluded that P could implicate the beliefs in A . For much the same reasons, noone of the a_i can be sufficient to believe any of the beliefs in the corresponding B_i . And again, for each b_{i_j} in B_i , the sufficient set cannot be found among the other beliefs in B_i . So we have to assume the existence of yet another set C_{i_j} for each b_{i_j} in B_i . And we have the very same problem. We must always avoid loops, because of the transitivity of the implication, so we must always assume new sets of beliefs.

Could these sets be made up of the same elements? Well, since a finite number of elements has only a finite number of combination, it would be impossible to avoid loops. So it seems that there must be an infinite number of beliefs in order to believe P and all the beliefs that make Smith in the condition of believing P . This is bad news for Fodor and Lepore, because, according to their definition of holism, it is only the number of beliefs that counts. According to our definition of holism, though, this is not yet a holistic picture, because we have not proved that all the belief system of Smith is involved in making him believe P . Smith's beliefs can form several different infinite

²⁸ *Ibid.*, pgs. 27 f., emphasis in text.

²⁹ Note, incidentally, that this is another proof that such a reading is not incompatible with Premise 1.

networks, unrelated one to the other.³⁰ Consider, however, that we have already proved that there is no way to draw a univocal distinction between beliefs in one network and beliefs in another. All we lack, then, is a proof of how the systematicity of the whole set of Smith's beliefs make it possible for him to believe *P*.

Actually, one problem with Argument A being considered as an argument for holism is that there is no mention of systematicity in it. This could be the reason why Fodor and Lepore had some problems in finding explicit instances of the argument in holistic authors.³¹ If I am correct, this is because the argument is missing something that holistic authors might find essential: mention of a system.

Fodor and Lepore's argument is claimed to have consequence also with regard to translation holism:

Translation holism wouldn't follow either, because it requires that for two people to share any belief, they must share at least another belief, and the present model allows that what you believe is *P* and *A*, whereas what I believe is *P* and *B*. What everybody *really* wants is that meaning should be anatomic and that translation holism should nevertheless be false. This suggestion seems to do the trick.³²

This means, more or less: "Holists tell us a proposition *P* is intelligible only with reference to its cognitive context; how do they deal with an argument showing that the same proposition *P* can be understood both in the context *A* and in the context *B*?" I think I can show there is no contradiction here. In fact, it can be seen as mere begging the question to individuate the same proposition *P* in two different contexts. How do we individuate *P* out of its relation with *A* or *B*? Until we are told that, we cannot evaluate properly the counterargument, and we have reasons to believe that holism would not accept identification of *P* independently of its context.

4.2 Translation and Interpretation

What we saw discussing Fodor and Lepore objections up to this point can be seen as a general view on holism. Differences about aspects of meaning have not yet proven relevant. Let us now begin our survey of interpretation, and see how a theory about it can be holistic. As I said, there is a sense in which every task of interpretation shares features with a translation from a language into another. No two persons, in fact, can be assumed to speak exactly the

³⁰I am indebted to Vincenzo De Risi for this point.

³¹*Ibid.*, pg. 26.

³²*Ibid.*, pg. 28, emphasis in text.

same language. Semantic holism (and perhaps semantic theories in general) cannot ignore differences between idiolects, because it cannot be decided in advance that there are irrelevant differences.

Translation has been a big issue for philosophers: to explain how languages with very different means are able to convey the same meanings – or whether they are – would indeed constitute a big step towards an explanation of what meaning is. To understand the terms of the question, we need to introduce some terminology, which I take from Donald Davidson:

In the general case, a theory of translation involves three languages: the object language, the subject language, and the metalanguage (the languages from and into which translation proceeds, and the language of the theory, which says what expressions of the subject language translate which expressions of the object language). And in this general case, we can know which sentences of the subject language translate which sentences of the object language without knowing what any of the the sentence of either language mean (in any sense, anyway, that would let someone who understood the theory interpret sentences of the object language).³³

The problem assumes its biggest proportions in the case of **radical translation**, i.e., in situations where object language and subject language are radically separate, in the sense that they are not akin, nor there are speakers someway competent in both languages:

Translation between kindred languages, e.g., Frisian and English, is aided by resemblance of cognate word forms. Translation between unrelated languages, e.g., Hungarian and English, may be aided by traditional equations that have evolved in step with a shared culture. What is relevant rather to our purposes is *radical* translation, i.e., translation of the language of a hitherto untouched people. The task is one that is not in practice undertaken in its extreme form, since a chain of interpreters of a sort can be recruited of marginal persons across the darkest archipelago. But the problem is the more nearly approximated the poorer the hints available from interpreters [...]. I shall imagine that all help of interpreters is excluded.³⁴

We will come to the conclusion that all interpretive tasks, even those which looks the simplest, are just as difficult to explain. This will be proved by means of an analysis of the radical task, and a few considerations about it.

³³Davidson (2001), pg. 129

³⁴Quine (1960), pg. 28, emphasis in text.

4.2.1 Some Issues about Behavioristical Theories of Interpretation

When Quine inquired the problem of radical translation, he wanted to survey, as if with a magnifying glass, an extreme case of a more general question – that of interpretation – so to prove the following general conclusion:

[...] the infinite totality of sentences of any given speaker's language can be so permuted, or mapped onto itself, that (a) the totality of the speaker's dispositions to verbal behavior remains invariant, and yet (b) the mapping is no mere correlation of sentences with *equivalent* sentences, in any plausible sense of equivalence however loose.³⁵

It is plain to see that this is a behavioristic conclusion. The point, according to Quine, is that we can identify the meaning of a sentence only through the behaviour of speakers, because we have no elements to treat meaning as an independent entity.³⁶ The meaning of a sentence, therefore, depends on the conditions which solicit uttering by the speaker. Yet, we have no guarantee that, for a given utterance, we can identify these conditions with certainty. Moreover, the same stimuli can solicit uttering of different non-equivalent sentences. That such a behavioristic theory about interpretation is a holistic theory of language is proved by what follows.

If meaning is to be recovered by means of behavioral survey, but behavior is always ambiguous, we have two possibilities. We may decide to attribute to the speaker a different meaning any single time he utters something. In other words, we may decide to interpret each utterance as if it were part of some new and until now unknown language – a language that will coincide in practice with the utterance itself, hence with the sentence which the utterance instantiates. The implication of this perspective is that language is neither unitary nor systematic. We could have been able to recover such features if we chose to investigate language from its semantic productivity, as atomistic semantics does; we decided, though, to go the other way round, and now we lack the evidence to assume systematicity and unity. As we will see soon, we are not at all forced to renounce them, but we cannot have atomism on one side and systematicity and unity on the other, if our approach starts from interpretation.

The other possibility is, hence, to renounce atomism to stick to a systematic and unitary picture of language. If we regard each utterance as carrying a meaning that we can recover only taking into consideration other utterances

³⁵ *Ibid.*, pg. 27, emphasis in text.

³⁶ As a theory of interpretation, this goes along the lines of what we have seen in § 2.2, with the severe restriction that a luxurious ontology is definitely unwanted. I recall that in § 2.3.1 we reached the conclusion that all interpretive processes are at least partly holistic, and that a minimal account of interpretation takes the form of behaviorism.

(possibly every other utterance produced by the speaker), then the speaker will obviously appear to be speaking a systematic and unitary language.

It is perhaps possible to claim that not every other utterance is needed: maybe only some of them – possibly only a few – would do. This would entail that a single language is divided in several sets of utterances. If we want language to be unitary and systematic, either we think that each of these sets of sentences is, in fact, a language, or we find some other way of dealing with this *molecularized* language.

For what concerns the latter position, i.e., what we called semantic molecularism, recall that, before we can confine our survey to just the some utterances needed to interpret the given one, we need a methodology to distinguish relevant utterances from irrelevant ones. It is undisputable that behavioral analysis, without further assumptions, cannot trace such a distinction.

For what concerns the former option, we should notice that here we would have the result that one single speaker speaks several languages, which is fine, as long as we are allowed to explore the semantics of each of these languages independently from the others. In what follows, therefore, I will keep considering each speaker as if speaking one single language.

Let us assume, thus, that to understand the meaning carried by an utterance, we must take all other utterances by the same speaker into consideration. Anyway, this procedure does not give us quite a good picture of language. As we have seen in § 1.2, the identification of a language with some set of utterances allegedly produced in that language clashes with the Humboldtian characteristic, i.e., with the apparent possibility of each language to produce infinite different sentences. Hence, we decided to talk of a language as a set of rules and a lexicon, which is a rather systematic and unitary view of language indeed. We must attribute it to a speaker as a whole, not piece by piece. Behaviour-based tests are, thus, tests to the whole linguistic system, not to parts of it.

The first thing we must do to interpret a sentence, according to the behavioristic theory seen so far, is to individuate the stimulus that prompts an utterance of it. This is equivalent, according to Quine, to individuate what makes the speaker deem the sentence true. It is clear that such a behavioristic theory has to make assumptions about truth and meaning much before having settled the rule system of the language. This is apparently a problem, but, in fact, we have seen in § 2.3.1 that every trial and error procedure involves holism, and all we need to survey the linguistic system is precisely a trial and error procedure. Holism is needed for such a methodology, because it is the only assumption that gives us freedom to look for relevant utterances throughout all the linguistic behaviour of the speaker, without biases that limit our survey.

4.2.2 Again on the Analytic/Synthetic Distinction

The conclusion that we reached in the previous section is acceptable only if we cannot assume that we can recover the meaning of sentences by appeal to analytic inferences and analytic truths permitted by the sentences themselves. Fodor and Lepore take Quine's "Two Dogmas of Empiricism" into consideration, to prove that no consequence as for semantic holism comes from the theory expressed in the paper.

Actually, they are skeptical about Quine having endorsed meaning holism in that paper: they deal with it because it has often been brought as a support to meaning holism. Indeed, Fodor and Lepore are right in thinking that the main argument of Quine's paper is what they call confirmation holism – i.e., the thesis that statements about the world can be evaluated only in connection with all the other statements of the theory they are part of. This is renown under the name of Duhem-Quine or Quine-Duhem thesis (Q/D thesis in Fodor and Lepore's abbreviation). One of the possible formulations of the thesis we have already met:

Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system. [...] Conversely, by the same token, no statement is immune to revision.³⁷

They may be right in thinking that Quine, at least in this paper, was not interested in semantic holism. However, I would be more cautious on this point, because Quine talked rather interchangeably about languages and about theories,³⁸ and this seems one of those cases where he claimed the distinction to be ininfluent.³⁹

However, for the present discussion, it is of no importance whether Quine was or not a holist about meaning, and whether he had or not an argument for being or not being one: as a matter of fact, he is believed to have been one. Here is the argument which semantic holists are claimed to have found in "Two Dogmas of Empiricism," according to Fodor and Lepore. The argument raises putting together confirmation holism and verificationism:

The Q/D thesis says that confirmation is holistic; that is, that every statement in a theory (partially) determines the level of confirmation of every other statement in the theory. Verificationism says that the meaning of a statement is determined by its confirmation relations. The invited holistic inference is that every statement in a theory partially determines the meaning of every other statement.⁴⁰

³⁷Quine (1953), pg. 43.

³⁸Cf. Chomsky (1969), pg. 53.

³⁹Cf. Quine (1969), pg. 310.

⁴⁰Fodor and Lepore (1992), pgs. 41 f.

What this argument really means depend on what we think is an appropriate definition of ‘statement,’ and this can depend, in its turn, in what we think that Quine was actually claiming when defending the Q/D thesis.

Statements

Fodor and Lepore claim that Quine’s preferred reading of the Q/D thesis (which we have just quoted) makes it necessary that we have some independent way of identifying statements across theories. Their proposal is to consider semantic properties as essential in individuating statements. Anyway, if we do, no holism can hold, neither about confirmation nor about meaning.

Confirmation holism would not follow because it is the thesis that confirmation relations are not essentially related to a given statement. That, together with the doctrine that confirmation relations have a semantic nature, entails that statement cannot have essential semantic properties.

Meaning holism, on the other hand, is the doctrine that sentences have a meaning because of the language they are sentences of, i.e., their semantic properties are derivative and not essential. Thus, meaning holism is contradictory with the posit of essential semantic properties of sentences.

So, our proposal would be of not identifying statements through semantic properties. Fodor and Lepore claim that it is doomed to fail, for two orders of reasons. First, we leave statements ambiguous for what concerns both their meaning and their linguistic affiliation. Second, this entails that we cannot be sure whether we have one or more statements in cases of ambiguities. I propose to solve both problems by stating clearer how to assign sentences to languages. Let us first briefly see how this would solve Fodor and Lepore’s problems, and then try to develop a theory to do so.

Semantic and Linguistic Contingency The first problem that Fodor and Lepore have with statements being identified by means of their morpho-syntactic properties⁴¹ is that not only the semantic properties of statements are contingent, but that also the “linguistic affiliations” of statements are contingent.⁴²

However, consider that a certain part of the semantic properties of every sentence is definitely dependent on its morpho-syntactic properties.⁴³ Thus, take Fodor and Lepore’s example: “The duck is ready to eat.”⁴⁴ If we individuate this statement by means of its morpho-syntactic properties, it cannot be ambiguous. For the syntactic properties of the statement are exactly what

⁴¹They do not make any mention of phonological properties, but I think that the argument is quite indifferent to their presence.

⁴²*Ibid.*, pg. 45.

⁴³Cf. Borer (2005a) and Borer (2005b).

⁴⁴Fodor and Lepore (1992), pg. 45.

make the difference between the reading in which “the duck” is the object of “eat” and that in which it is its subject.

Also, the morpho-syntactic properties of a sentence are dependent on its linguistic affiliation, i.e., from what language it is a well-formed formula of. Vice versa, the morpho-syntactic properties of a sentence may help us in defining its linguistic affiliation. Even if some uncertainty among all the different idiolects that allow it as a well-formed formula will remain, it does not follow that the statement is contingently a statement of one or the other idiolect: on the contrary, given the morpho-syntactic properties of the statement and the rules of those idiolects, it is necessary that the statement is allowed as a well-formed formula by any of those idiolects.

In claiming semantic and linguistic contingency, Fodor and Lepore seem to be speaking of statements as utterances, which have, indeed, “both their semantic properties and their linguistic affiliations contingently.”⁴⁵ This is because utterances, taken as such, lack morpho-syntactic properties by which we could individuate them and which give them some semantic properties and linguistic affiliation.

Note, also, that, if semantic holism is true, linguistic affiliation and semantic properties are not to be distinguished: the semantic properties depend on the linguistic affiliation. And holism accounts for the vaguenesses and ambiguities we have found: a statement in isolation, taken with its morpho-syntactic properties cannot be assigned to exactly one language and we cannot know its exact meaning because to do so we need the other statements of the language. Or such is the holistic claim.

Now, it is true that morpho-syntactic properties are not enough to make it possible for us to assign a meaning and a language to a statement. They are enough, though, to make a rough assignment that can be exploited as a base for further investigation, and that could be subject to revision given enough data about other statements. All this issue will become clearer when I will enter in more detail about the process here involved, which I have to postpone until § 4.2.7.

Semantic Individuation Fodor and Lepore have problems with the morpho-syntactic individuation of statements also because they think that it is a truism that statements are individuated by their semantic properties.⁴⁶ My first answer to this is that it cannot be a truism at all, since we do not know yet what statements are.

Fodor and Lepore think they need such an assumption to build the following argument, “in the spirit of ‘Two Dogmas,’”⁴⁷ to show that a system (be it a theory or a language) could not contain only one statement, e.g., the

⁴⁵ *Ibid.*

⁴⁶ *Ibid.*, pg. 44.

⁴⁷ *Ibid.*

statement that it is raining:

Premise 1: The statement that it's raining (R) is partially confirmed by the statement that the streets are wet (S). (Metereological platitude.)

Premise 2: Confirmation relations are ipso facto semantic. (Peirce's thesis [i.e., verificationism].)

Premise 3: Statements are individuated by their semantic properties; or, as we will sometimes say, they have their semantic properties essentially. (Truism.)

Lemma: R is individuated, inter alia, by its relation to S.

Conclusion: Any theory that contains R must contain S. A fortiori, no theory could contain just R.⁴⁸

The problem is with Premise 3, of course. It does not seem to me that such a premise is very much in the spirit of "Two Dogmas." For Quine believed that "the unit of empirical significance is the whole of science."⁴⁹ Fodor and Lepore claim that this formulation is ambiguous,⁵⁰ and we can agree with them. But something must be clear in spite of all the ambiguity. If, according to verificationism, confirmation relations are ipso facto semantic, then statements have no semantic properties deriving directly from confirmation relations. Because, if there is a non-ambiguous reading of Quine's claim, it is that only the whole science is susceptible of being confirmed or disconfirmed. Therefore, statements cannot have their semantic properties essentially, but only indirectly, thanks to the role they play inside a theory.

Now consider how this translates when we are talking about language. It seems that all boils down to the platitude that a sentence can have a certain meaning only if considered as a sentence of a certain language. And now we must ask how we decide which language a sentence is formulated in.

Take (1), for instance. It is an ambiguous utterance, inasmuch as it can mean "Romans' calves are beautiful" in Italian and "Go, o Vitellius, to the sound of war of the Roman god" in Latin.

(1) I Vitelli dei Romani sono belli.

Do we have one sentence with two possible meanings or two different sentences? Fodor and Lepore would say that, if meaning comes not into the question as an essential property of sentences, we are presented the same sentence. I would say that, since sentences are related to languages and languages to speakers, then the decision about meaning and language should take speakers into consideration.

⁴⁸ *Ibid.*

⁴⁹ Quine (1953), pg. 42.

⁵⁰ Fodor and Lepore (1992), pgs. 40 f.

Now, one difference between languages and theories is that a theory is essentially shared, while a language can be thought of as essentially individual. The notion of linguistic community is vague and perhaps derivative, for our purposes, of that of idiolect.⁵¹

On this regard, I propose to compare how we assign statements to theories and how we compare theories with the case of languages. The notion of common language, as I said, is vague, as it is often difficult to state which idiolect is part of which common language, and, therefore, what are the characteristics of common languages. Theories, when we talk about them seriously,⁵² have no such vagueness. In some respect, public scientific theories will show to be like idiolects.

Is the Q/D Thesis Trivial?

Back to the general case of Q/D thesis, then. Suppose we decide to hold onto statement *s*, come what may, and that to do so we have to make some adjustments to other statements in our theory. If the meaning of *s* did not change after changing the theory, holism would not hold. If it changed, though, Fodor and Lepore would claim that we have trivialized the Q/D thesis:

It's only epistemologically interesting that you can hold onto "Burning is the liberation of phlogiston" even in the face of Lavoisier's results if "Burning is the liberation of phlogiston" means that burning is the liberation of phlogiston. It's no news that you could hold onto it in the face of those results if it means that Greycat has whiskers.⁵³

I propose to show that this trivialization can be avoided without assigning unchangeable meanings to the statements. The idea, as I have already suggested, is to pay attention to what it is for scientific statements to be true or false in the context of a theory. I think, as I have said, that this can give us some insight in what it means for a sentence to have a meaning in a language.

This Way Out The way to save Quine's reading of the Q/D thesis, together with the holistic desire to have statements change their meanings depending on the theory they are in, can be to exploit an idea by Kuhn:

The claim that two theories are incommensurable is [...] the claim that there is no language, neutral or otherwise, into which both

⁵¹Cf. § 1.2.

⁵²That is, as "fairly specific bod[ies] of speculation or propositions with a definite subject matter" (Hacking (1983), pg. 175). I am not claiming that this is the right interpretation to give to Quine's discussion of theories, but that this interpretation is consistent with what he said, and, for independent reasons, the most convenient (cf. *ibid.*, pgs. 174 ff).

⁵³Fodor and Lepore (1992), pg. 47.

theories, conceived as sets of sentences, can be translated without residue or loss. [...] Most of the terms common to the two theories function the same way in both; their meanings, whatever those may be, are preserved; their translation is simply homophonic. Only for a small subgroup of (usually interdefined) terms and for sentences containing them do problems of translatability arise.⁵⁴

To understand what use I have in mind for this idea, consider this passage from one of Fodor and Lepore's notes:

[...] Quine likes an ecumenical story according to which the connectives *mean different things* in classical and intuitionistic logic. But, again, barring some trans-theoretic notion of statement identity, it's unclear how you decide *which* connectives it is that the two kinds of logic assign different meanings to. (The thought that there are perhaps *no* theorems that classicists and intuitionists can both accept seems, to put it mildly, unintuitive).⁵⁵

On this regard, I have to acknowledge the insuperable chasm between my notion of what is intuitive and Fodor and Lepore's. Never would I have imagined that one could have intuitions about theorems of formal logic.

If intuitions were not an issue, though, it would seem obvious to me that Quine was referring to a way of identifying connectives by way of their symbolic role in the theory. When we compare theories, we do so under the point of view of some theory. Correspondingly, when we compare languages, we do so in a language. Therefore, we can decide to proceed to identify statements (hence, in the present case, connectives) in a homophonic way, and acknowledge the shift of meaning by their presence in theorems that are accepted in one formal theory but whose homophonic counterparts are not accepted in the other.

As in the case of classical connectives vs. intuitionistic ones the shift that meaning is allowed to do intervenes on a preceding homophonic identification of statements and connectives, so, in the former example, the shift that we must do in the meaning of, say, "phlogiston" (or, maybe, of the phrase "liberation of phlogiston") in order to be able to hold onto "Burning is the liberation of phlogiston" even in the face of Lavoisier's results is limited, because it cannot be arbitrary. It is precisely the shift that is triggered by the fact that, now, "phlogiston" will appear in some true statement that used to be considered false. Fodor and Lepore could ask what determines the meaning of these statements, then. I answer that meaning assignation is achieved at first by homophonic connections, just as Kuhn believed, and then by subsequent adjustments and feedbacks from one statement to the other.

⁵⁴Kuhn (2000), pg. 36. Cf. also § 6.2.1 below.

⁵⁵Fodor and Lepore (1992), pg. 218 n. 10, emphasis in text.

Sameness of Meaning? Note that we do not have said much when we say that we want “Burning is the liberation of phlogiston” to mean that burning is the liberation of phlogiston. For the quoted statement is something that we want to be true (indeed, we want to hold onto it whatever the evidence may be), while, in our understanding of it, it is false that burning is the liberation of phlogiston – and this is because there is no phlogiston. Now, that seems to me like trivializing the Q/D thesis. I will change the example to get the point clearer. In the Middle Age they used to hold (2) as a true statement.

(2) Unicorns are white.

Suppose that, in the face of zoological evidence that there is no unicorn, we want to hold onto (2), and we want it to mean that unicorns are white. We want, i.e., (2) to have the same meaning of a proposition that assigns a colour to hypothetical, non-existent creature. Now, *that* is definitely epistemologically uninteresting.

My point is that we do not understand quite well what the meaning of sentences containing “phlogiston” or “unicorn” should be if they must have existing references. Of course, we can say that, whatever that meaning is, is different from that of (3). But this is not because we know that the meaning of (2) is different from (3), but only because we want it to be different and we stipulate it to be so.

(3) Greycat has whiskers.

If I am on the right track, this would imply that we cannot choose to hold onto one statement in isolation, but we must also decide to preserve some of its relations with other statements. This boils down to the unsurprising thesis that when we need to change our theory, it will not help to simply decide what we want to leave unchanged, but also what changes to operate.

Once again, what has all of this to do with language? Perhaps only confusion will ever arise from confounding language and theory. But maybe we can learn from the case we have raised that what we want to do with the linguistic side of the problem (i.e., finding a way to assign sentences to languages without assigning them essential meanings) could be done in a dynamic process, a process starting with assumptions, making hypotheses, acquiring data, revisiting and revising old hypotheses and making new ones, and so on. This is the kind of theory we will try to sketch in the next sections.

4.2.3 Semantics As an Empirical Enterprise

A clear consequence of the theory seen so far is that it leaves interpretation largely indeterminate. In fact, we cannot know from mere behaviour how to interpret others’ utterances, because we cannot be sure about which utterances

in our language should be considered their right translation. Quine proposed to introduce a principle of charity to reduce indeterminateness. He proposed, that is, to attribute to the speaker whose sentences we want to interpret a system of beliefs similar to ours. Of course, this will not always be possible:

Tolerance is bound to have been exercised if a native sentence, believed by the whole community with a firmness that no stimulus pattern of reasonable duration would suffice to shake, is translated as ‘All rabbits are men reincarnate’. To translate a stimulus-analytic native sentence thus into an English sentence that is not stimulus-analytic is to invoke translator’s licence. I think this account gives such a translation quite the proper air: that of a bold departure, to be adopted only if its avoidance would seem to call for much more complicated analytical hypotheses. For certainly, the more absurd or exotic the beliefs imputed to a people, the more suspicious we are entitled to be of the translations; the myth of the prelogical people marks only the extreme. For translation theory, banal messages are the breath of life.⁵⁶

In Quine’s system, such a principle of charity is just a prudence principle. Quine took evident advantage of it, but he gave no principled justification for it. In fact, the principle of charity seems to have at least partially the purpose of taking account of our intuitions; nevertheless, according to Quine, intuitions are legitimate in their territory, but we must sometimes put them aside when developing a coherent theory.⁵⁷

We may decide to give to the principle of charity a more central position in our theory of interpretation.⁵⁸ If we decide to renounce an utterly behav-

⁵⁶Quine (1960), pg. 69. Stimulus is, in Quine’s theory, a time-dependent variable. Those sentences that would solicit assent no matter what stimulus the speaker is faced with are considered analytic. But it suffices to vary the temporal gap considered to let the distinction fall. For instance, if we considered as stimulus a speaker’s entire life, the category ‘analytic’ would doubtlessly lose any point. Cf. *ibid.*, pg. 63.

Quine spoke of “stimulus-analytic” sentences since he did not believe in the possibility of drawing distinctions between synthetic and analytic sentences not grounded on behaviorist bases. Cf. Quine (1953), chpt. 2.

⁵⁷Cf. Quine (1960), pgs. 66 f., where the issue is with analyticity, not with the principle of charity:

The intuitions are blameless in their way, but it would be a mistake to look to them for a sweeping epistemological dichotomy between analytic truths as by-products of language and synthetic truths as reports on the world. I suspect that the notion of such a dichotomy only encourages confused impressions of how language relates to the world.

Without a theory to support them, intuitions can at best furnish us a guide in our investigation, but never constitute integral part of it. I think that there is no reason to limit these considerations of the drawbacks of resorts to intuitions to those about analyticity.

⁵⁸Cf. Davidson (2001), pgs. 228 f.

ioristic account, we can give intuition and introspection the place they seem to have in our daily practice of linguistic interpretation.⁵⁹

According to Quine, we have seen it, interpretation is an empirical enterprise; therefore, it needs empirical hypotheses. More than one hypothesis can be consistent with the same evidence. To see how a principle of charity can help in individuating the right theory, we must first state what we expect from an empirical theory of interpretation.

Let us begin with Tarski's definition of truth, as we saw it in § 1.3. It consisted in defining truth for a formal language \mathcal{L} ('object language') with respect to another language \mathcal{M} ('metalinguage'), containing a predicate 'true in \mathcal{L} ' which cannot be a predicate of \mathcal{L} itself. If \mathcal{M} contains a constant for every sentence p of \mathcal{L} , and, for every sentence p_n of \mathcal{L} , a sentence s_n translating it, we can say, in \mathcal{M} ,

‘ p_n ’ is true in \mathcal{L} if and only if s_n .

Call sentences of this kind T-sentences. Tarski could construct meanings by other means, and his is only a theory of truth. As such, though, it tells us nothing about meaning. In other words, we cannot exploit the correspondence created by T-sentences to ground the notion of meaning, and to find what sentence in \mathcal{L} mean.

Donald Davidson proposed another way of exploiting T-sentences. He suggested that any acceptable theory of meaning should entail all and only the true T-sentences. T-sentences, thus, should be used as a test to compare inferences of a theory of meaning developed by other means, rather than as the starting point of a theory of meaning. This enables us to see semantics as an empirical enterprise: we know, in fact, what results we expect the theory to produce as its outcome, and, therefore, we can test it.

Davidson's problem is to explain how the right kind of theory should be developed. Consider that we are looking for a theory which entails all the true T-sentences like “‘Snow is white’ is true in English iff snow is white,” while forbidding T*-sentences like “‘Snow is white’ is true in English iff grass is green.” Note that the problem with T*-sentences is not that they are false, therefore we must find another criterion to discriminate between them and actual T-sentences. Call this “the extensionality problem.”⁶⁰

Another requirement for our theory of meaning that we must consider is that it must start abruptly without former knowledge of meaning in \mathcal{L} . Otherwise, we should ask where that knowledge comes from. In other words, we are asking for a theory of meaning that is achievable by a radical interpreter, someone that is faced with the problem of radical interpretation.

⁵⁹We might notice how, from the theories seen in § 2.2, to the behaviorism just delineated, to the theory we are about to describe, we are increasing the importance of considering language as a human faculty and activity, and not just a theoretical independent entity.

⁶⁰Cf. Fodor and Lepore (1992), pgs. 60 ff.

To summarize, Davidson problem was to solve the extensionality problem, while meeting the radical interpretation requirement. This matters to us, because it can be seen as a reformulation of the problem we had in last section. There we wanted to know the meaning of sentences (and this we can do by solving the extensionality problem) within a theory which does not individuate sentences via their meanings (that is, a theory within which sentences are initially taken as susceptible of any meaning, i.e., a theory that meets the radical interpretation requirement). Davidson's problem is just our problem.

Fodor and Lepore individuate and discuss three possible solutions to the extensionality problem suggested by Davidson:

Suggestion 1: Exploit the fact that natural languages exhibit compositional semantic structure; in particular, that the same expressions can recur with the same meanings in (indefinitely) many formulas. This is the main strategy contemplated in "Truth and meaning."

Suggestion 2: Require that the T-sentences in favored truth theories be laws. This is the strategy contemplated in the "Reply to Foster" and in note 11, added to "Truth and meaning" in 1982.

Suggestion 3: Require the favored truth theory for L to entail T-sentences according to which most of the sentences that speaker of L hold true *are* true. This is one version of Davidson's "principle of charity" and is the main strategy contemplated in "Radical interpretation."⁶¹

This is the order in which Fodor and Lepore discuss the suggestions, and they interpolate a discussion about the plausibility of the radical interpretation requirement between the first and the second. Finally, they briefly dismiss the question about all three suggestions working together. Maybe they should have paid more attention to this final complex suggestion, because if it is true, as they claim, that "the issues raised by suggestions 2 and 3, unlike those raised by suggestion 1, are intimately concerned with the question of radical interpretation," then suggestion 1 should always be considered in connection with the other two.⁶² For Davidson held the radical interpretation requirement to be unavoidable. Anyway, we will discuss this at due time.

⁶¹Fodor and Lepore (1992), pg. 62, emphasis in text.

⁶²This is confirmed, for instance, by the fact that in "Truth and meaning" there is mention of radical interpretation and of the principle of charity (Davidson (2001), pg. 27). Note 11 is claimed by Davidson to be a clearer restatement of the text it refers to, but this should be taken *cum grano salis*.

4.2.4 Compositionality

The first suggestion considered by Fodor and Lepore is that our theory of meaning derives T-sentences “by exploiting the linguistic structure of the formulas whose truth conditions they specify.”⁶³ The theory, therefore, must derive the semantic properties of one sentence from the semantic properties of its lexical parts considered together with its syntactic structure. This constraint is holistic in the sense that the theory of meaning that we must derive should hold for all sentences at the same time, hence we must avoid contradictions between the semantic properties of lexical items appearing in different sentences, or between syntactic structures shared by different sentences. Thus, e.g., a T*-sentence like “‘Snow is white’ is true in English iff grass is green” “could not belong to any reasonably simple theory that also gave the right truth conditions for ‘That is snow’ and ‘This is white.’”⁶⁴

Necessity of Compositionality

Fodor and Lepore comment that this suggestion entails that there cannot exist noncompositional languages, for the expressions of such languages could not have determinate truth conditions.⁶⁵ But, in fact, such a language is not impossible, or so they claim:

Suppose there is a child who has (to all appearances) mastered the entire nonrecursive apparatus of English. So he can say things like “It’s raining,” “Snow is white,” “Grass is green,” “That’s snow,” “That’s green,” “That’s frozen,” “Everybody hates me,” “I hate spinach,” and so forth, but not “Snow is white and grass is green” or “Everybody hates frozen spinach,” and so forth. We take it that, assuming that everything about the child’s speech dispositions in respect to this nonrecursive part of the language is exactly like that of the corresponding normal adult, it is *very* plausible that when the child says “Snow is white,” he means that snow is white. [...]

But now consider a child who is *just like* this one in his speech (and inferential and, generally, cognitive) dispositions except that, whenever child 1 would use “Snow is white” to say that snow is white, child 2 uses the unstructured expression “Alfred”; [...] and so forth.

[...] Then, barring tasks that require the comprehension or production of recursive sentences, let it be that the translated verbalizations of child 2 are *indistinguishable* from the verbalizations

⁶³*Ibid.*, pg. 63.

⁶⁴Davidson (2001), pg. 26 n. 10.

⁶⁵Fodor and Lepore (1992), pg. 65.

of child 1. *Nevertheless*, if compositionality is a necessary condition for content, then there is an a priori argument that child 2 couldn't mean anything determinate by what he utters. We take it that is just obvious that such an argument couldn't be sound. After all, whether the child means anything by his utterances probably depends on the intentions with which he utters them. What a priori argument would show that a child couldn't utter "Sam" with the intention of thereby saying that snow is white?⁶⁶

I think that there are several problems with this thought experiment. First of all, it is not clear what the appeal to intentions should solve. In fact, it is yet to be proved that the meaning of sentences is dependent on the communicative intentions of the speaker. Even if this is so, our attribution of intentions to child 2 seems to be dependent on some kind of intuition, and it will become clear that nothing different from the principle of charity is at stake here. We will discuss it after introducing the theme of radical interpretation.

We can also ground intention attributions on behaviour. However, it is easy to see that child 2 has a linguistic behaviour very different from that that would allow us to attribute him any intention to mean, say, that snow is white when he utters "Sam." In fact, as Chomsky once pointed out, we do not know how to identify stimuli prompting such an intention.⁶⁷ Moreover, part of the intention to mean that snow is white is probably the intention to predicate something of something else. But in "Sam" there is no trace of predication, because it is unstructured.

On the other hand, we might say that intention attributions do not matter here, because intentions can ground meaning even if we are not able to identify them. This can be perfectly right. My only concern is why we would want to speak of language in such a case, given that we could not even trace a systematic and consistent use of utterances. The language spoken by child 2 would be the classical example of a private language, and there is something, in a private language, that strikes us as not linguistic at all. In fact, sentences in a language have a meaning even independently of intentions: for instance, I could intend to say that snow is white, but a slip of the tongue makes me say "Snow is wise." Does my utterance mean that snow is white or that snow is wise? I think that it means the second, and that is why I can say I was mistaken in uttering it. If this is true, then appealing to our intuition about intentions is pointless here, because intentions are not grounding meaning.⁶⁸

⁶⁶*Ibid.*, pgs. 65 f., emphasis in text.

⁶⁷Chomsky (1969), pgs. 57 f.

⁶⁸I incidentally underline that I am not saying that, in the example above, we must follow our intuition that "Snow is wise" mean literally. Rather, this is what our behavior towards such an utterance shows: we apologize and correct ourselves, state what we actually meant, and so on; but we usually do not claim that, in that kind of occasion, "Snow is wise" means "Snow is white."

It might turn out that something else is grounding the meaning of child 2's utterances. However, this does not seem the case. For consider that child 2 himself cannot know when he is mistaken. When he utters "Sam," he can never know whether, for some reason, "Mary" would have been more appropriate in that occasion. For that would mean not only that he knows the meaning of his utterances independently of his intentions, and we can grant it for the sake of the argument, but also that he is capable of knowing the propositional content of his intentions in order to compare it with his utterances' meaning. But in order to make such a comparison, he has to attribute propositional attitudes to himself, or else exploit some compositional device in order to attribute to the same entity both that it is an intention and that its content is "Sam." But our hypothesis was that child 2 has no compositional structure. Therefore, he cannot even know what his utterances mean, thus leaving us with the suspect that his might even be not a language at all.

Note that there would seem to be, theoretically, the possibility that child 2 does not actually know the meaning of his utterances, but that some (atomistic) symbol/world relation grounds meaning nonetheless. However, consider that if child 2 does not know anything about this symbol/world relation, he cannot use his symbols consistently with it. To postulate such a relation is pointless, because this relation has no effect whatsoever on the child, nor on his use of language, nor on the linguistic output itself, nor on the listener or interpreter. Only an omniscient being could know the real meaning of such utterances, and whether child 2 has intentions that correspond to that meaning:

Omniscience can obviously afford more bizarre theories of meaning than ignorance; but then, omniscience has less need of communication.⁶⁹

Consider child 1, now. If a case can be done about child 2, and we grant Fodor and Lepore that the cognitive dispositions of the two children are actually the same, this means that also child 1 is not really a speaker of any language. Indeed, his sentences only seem to have structure, because we can interpret them as structured, being able to connect them (homophonically) to sentences in our language, in which we can individuate syntactic and semantic components. But child 1 cannot do anything of the like, therefore he cannot compare "Snow is white" and "This is snow" as both being about snow. To do this, he should know that these sentences are compositionally structured, which he does not know for hypothesis. His sentences, then, are as unstructured as child 2's, and his language is as impossible as child 2's.

Note that we have to conclude that both languages have no semantics at all, and, therefore, they are not properly languages at all. Actually, we could assume, as noted before, that they have semantic properties, but that they are

⁶⁹Davidson (2001), pgs. 26 f.

simply inaccessible except for some omniscient being. Such an assumption, though, is in contrast with Davidson's requirement that a theory of meaning should be empirically testable.

What we have said about Fodor and Lepore's thought experiment can be used as a proof against the possibility of noncompositional languages. Consider a possible objection:

Children actually do go through a "holophrastic" stage in the course of normal language acquisition. Could there be a transcendental argument showing that the utterances they produce at this stage have no truth conditions? But, in fact, the child's holophrastic utterances are quite often *interpretable*.⁷⁰

The fact that something is interpretable, however, tells us nothing about its truth conditions and about the speaker's linguistic knowledge. Even dogs' whines are interpretable, and much involuntary behaviour as well is often considered interpretable. We can ascribe meaning to a sentence by a speaker, only if we assume the speaker has beliefs and desires similar to ours. But these beliefs and desires can be assumed only in presence of a linguistic behaviour mostly coincident with ours. There is nothing wrong in thinking that a child cannot have the semantic subtleties of an adult. And there is nothing wrong, thus, in thinking that a child does not mean what an adult means in uttering utterances homophonic to the adults' ones. Of course, this line of argument is leading us towards the issue of interpretation. This will be the topic of the sections from 4.2.5 on.

Sufficiency of Compositionality

Before analysing the role of radical interpretation in Davidson's theory of meaning, we need to deal with another piece of argument by Fodor and Lepore. Not only they claim that compositionality is not necessary to solve the problem of extensionality, but that it is also not sufficient. In fact, if we consider a pair of coextensive but not synonymous atomic predicates, then there will be no compositional test that will help us distinguish them, unless we appeal to intensional contexts.⁷¹ This seems to have unwanted consequences:

But if the compositionality solution to the extensionality problem must appeal to *these* contexts, then it seems to be implied that *there could not be an entirely extensional language whose sentences have determinate truth conditions* (contrary, it goes without saying, to the widely held view that, really, there couldn't be any other language whose sentences do).⁷²

⁷⁰Fodor and Lepore (1992), pg. 224 n. 6, emphasis in text.

⁷¹*Ibid.*, pgs. 66 f.

⁷²*Ibid.*, pg. 67, emphasis in text.

Such a consequence does not seem to be really entailed. In fact, let us ask how we identify coextensive but not synonymous atomic predicates. Again, it seems that some sort of omniscience is required. But the very assumption that such a pair can exist in a language lacking the means to distinguish them (such as a language deprived of intensional apparatus) seems to beg the question. Again, in fact, an atomistic symbol/world relation is assumed, and such relation cannot be described in any way, because it bears no relation to the way language is used or to the way we interpret it. Does such a notion of meaning have any sense or utility?

Moreover, consider that an atomistic theory of meaning would not be able either to deal with the extensionality problem without appealing to intensional contexts and semantic properties of such contexts. So, the problem spotted by Fodor and Lepore is a problem for all theories of meaning, and cannot therefore be used to discriminate among them.

Our solution to this problem, then, will be simply to bite the bullet, but not in the sense wanted by Fodor and Lepore. We will not claim that there is no successful theory of meaning for extensional languages. We will state, instead, that in the case of this languages it is pointless to speak of pairs of coextensive though not synonymous atomic predicates. In an extensional language all coextensive atomic predicates are synonymous.

If this is true, then the other objection to the appeal to compositionality advanced by Fodor and Lepore fails too. They claim that the term “intensionality” is not at a semantic metatheory’s disposal. In fact, there is no way for such a theory to identify intensional contexts before the theory of meaning is complete. I am perfectly fine with such an objection, provided that from the present discussion should be clear that no mention of intensionality is needed at all. We can do exactly how Fodor and Lepore, as it seems, assume: after having provided a theory of meaning and having tested it, we recognize coextensive atomic predicates and we decide whether they are synonymous or not by looking for contexts discriminating among them. Contexts that will do will be called intensional. If no such context can be found, then we will assume the predicates to be synonymous. This, in fact, is the way they are used, and this is the way the speakers themselves interpret them. Any further consideration would imply semantic omniscience.

On this regard, let me stress again the difference between the two approaches to semantics that here are at stake. Our discussion of semantic theories shows that there are two ways of dealing with meaning in natural languages. One way necessarily take into consideration only empirically testable hypotheses. Therefore, we will end necessarily with a theory that makes assumptions on the nature of language on the basis of the evidence that we have for it. Nevertheless, it could be that the real metaphysic nature of meaning does not, in fact, mirror the phenomenic nature of meaning, which we study. From this point of view, any theory of meaning can be right or wrong. However,

right and wrong themselves seem inappropriate in such a context, because we cannot ever decide which is the right theory and, moreover, even if we knew, it would not help us in everyday use and study of language. Therefore, the way that starts from interpretation of language and leads to a corresponding theory of meaning seems the only tenable. And this seems the reason behind Davidson's requirement that theories of meaning must be empirically tested.

The empiricist requirement can also account for another objection put forward by Fodor and Lepore. They note that if a theory entails the T-sentence “‘Snow is white’ is true in English iff snow is white,” and LT is some accepted logical truth, then the theory entails also “‘Snow is white’ iff snow is white and LT.” Davidson's response to that is that a derivation of T-sentences entailed by a theory of meaning should be canonical, in the sense that no additional axioms should be required. The aberrant T-sentence of our example, then, should be ruled out, because it requires the addition of logical apparatus to the core of the theory of meaning.⁷³

Fodor and Lepore reply that the aberrant T-sentence could be derived not only by appending “and LT” to the canonical T-sentence, but also from an axiom added directly to the core set of axioms of the theory of meaning. I think that we can respond along Davidson's lines that, since the theory of meaning must be an effective empirical theory, we must have justifications for adding axioms.

Punctate Languages

Fodor and Lepore end their discussion of the compositional solution stating that, even if the problems they found should be mended, you could nevertheless have a language all of whose sentences are token-reflexive:

[...] there is no argument so far that you can't have a language *all* of whose sentences are token-reflexive. [...] If, however, you can have a language that contains only token-reflexive sentences, what argument shows that you can't have a language that contains only one sentence, so long as *it* is token-reflexive?⁷⁴

The problem with this objection is again that it overlooks the empirical nature of a theory of meaning. In fact, how can we understand what an utterance of the intended kind (e.g., “This is blue”) means? We cannot test our hypotheses: we cannot put blue objects in front of speakers and expect them to utter, “This is blue.” Intentions are beyond our survey, because we need other linguistic clues to decipher them. Maybe we can utter the sentence with an interrogative inflexion? But interrogative sentences are different from

⁷³*Ibid.*, pgs. 67 ff.

⁷⁴Fodor and Lepore (1992), pgs. 69 f., emphasis in text; cf. § 2.1.2.

affirmative ones, therefore they are by hypothesis excluded from the language we are questioning. Worse than that, speakers themselves cannot interpret their own language, wonder about the meaning of their utterances, nor claim, “‘This is blue’ is true in English iff this is blue,” because this is a different sentence from the initial one.

Perhaps someone [...] will be tempted to say, ‘But at least the speaker knows what he is referring to.’ One should stand firm against this thought. The semantic features of language are public features. What no one can, in the nature of the case, figure out from the totality of the relevant evidence cannot be part of meaning. And since every speaker must, in some dim sense at least, know this, he cannot even intend to use his words with a unique reference, for he knows that there is no way for his words to convey this reference to another.⁷⁵

4.2.5 Radical Interpretation

The discussion in the last section should have made it clear that what we think about the nature of meaning (i.e., a theory of meaning) has important connections with the way we come to determine meaning (i.e., a theory of interpretation). Of course the two theories are not the same one, because in the case of interpretation we can always imagine some cognitive device to cut corners, so to say. But they are not separable under two respects:

1. The two theories must be compatible. The explanation provided by one on its subject matter cannot contradict or entail consequences that contradict explanations and entailments of the other. Of course, this leaves plenty of space for accommodation and explanation of apparent contradictions.
2. A theory of interpretation furnishes the privileged way to test a theory of meaning, because the former is nothing more than a description of the actual empirical results of the axioms and theorems of the latter.

This is why an examination of the radical interpretation requirement is now needed. In particular, we must see if something similar to radical interpretation is ever possible, and, if it is, what the exact requirements for it are, and what the consequences of such requirements are for the holistic nature of the theory of meaning.

Viewed in this context, a radical interpretation (RI) theory seeks to do two things. First, it must specify the kinds of empirical evidence that a successful meaning theory may be required to account

⁷⁵Davidson (2001), pg. 235.

for. In this context, “empirical evidence” includes *any contingent propositions that the radical interpreter may legitimately appeal to to warrant his interpretation*. [...] Second, since different ways of constraining the evidence will lead to different T-theories being selected, an RI theory must *justify* the imposition of one set of evidential constraints in preference to others.⁷⁶

The first objection to RI theory is that it seems to contradict the Q/D thesis.⁷⁷ This is because “the justification of the evidential constraints that an RI theory imposes must be a priori,”⁷⁸ while “Q/D asserts that what counts as evidence for an empirical theory is always to be determined a posteriori.”⁷⁹

In fact, RI requires only that it is a priori determined what sort of evidence must be at disposal of the theory, but it does not determine a priori what pieces of evidence actually there must be. In other words, it simply states a priori how the evidence that we decide to treat as such should be considered, but it does not say what is considered as evidence given a certain theory. Thus, there seems to be no contradiction. In fact, even according to the Q/D thesis choices among the theories and among evidence are not always made a posteriori, as Fodor and Lepore claim. All sorts of requirements can be imposed on theories and on evidence, such as methodological constraints, ethic or esthetic considerations, metaphysical claims, and so on.

Talking of the possibility of RI, Fodor and Lepore identify two possible lines of argument. The first is to claim that some actual activities, like the work of field linguists, or the acquisition of language by children, have so much in common with RI that their possibility is direct evidence for the possibility of RI.⁸⁰

In both cases, Fodor and Lepore individuate some element that seems to get out of the RI requirements. For field linguists have of course a preexistent theory about language, which they exploit to discriminate among evidence, and children are claimed to have something analogous, a universal grammar (UG) that enables them to choose the right evidence and to organize that evidence in the right kind of theory.

Again, the contradiction seems to me only apparent, and it originates, probably, from not considering what the reasonable requirements for an RI theory should be, if we want it to be some empirical theory. Of course linguists have some preexistent theory, but either this is warranted or not. If it is warranted, for instance by experience, we should ask how this warranty was

⁷⁶Fodor and Lepore (1992), pg. 71, emphasis in text.

⁷⁷*Ibid.*, pg. 225 n. 12.

⁷⁸*Ibid.*, pg. 71.

⁷⁹*Ibid.*, pg. 225 n. 12.

⁸⁰*Ibid.*, pg. 73.

possible. We probably would end up with the very RI that we were looking for.

However, it seems more plausible to think that no empirical theory is completely warranted by experience, and this is essentially the Q/D thesis. Therefore, it seems that RI theories are not in a different position from sophisticated linguistic theories: both kind of theories are somewhat underdetermined by empirical data. This is not a problem: all that it shows is, as we already knew, that as in the case of the field linguists, the radical interpreter has to make some assumptions a priori, which will constrain his theory. It is not obvious that such assumptions should concern language. In fact, as we will see, the most important assumption of this bundle is the principle of charity, which is a methodological assumptions regarding, in its most general form, rationality as a human characteristic. Other methodological assumptions are certainly part of the bundle: assumptions regarding economicity of the theory, or elegance, and so on, which are shared by virtually all empirical theories.

As for the child's case, even granted that UG constitutes, indeed, a pre-existent linguistic cognition, again there is nothing particular different between the child's situation and the radical interpreter's. The major difference is that the child's initial assumptions (i.e., UG) are not conscious, while the radical interpreter's one probably are, being characterized as methodological requirements. This would account for the greater ability of the child in retrieving what we can call the right theory of meaning for the language. Consider, however, that this is a rather misleading way to speak: not much because it can be doubted that what the child has is, in fact, a theory,⁸¹ but because, in this case, there is no such a thing as the right theory. We must acknowledge that starting from the same data more than one theory is possible, and the only criterion to rule out the wrong ones is by way of T-sentences. Children are more efficient in doing so, probably because the conscious application of a rule can always be seen as a more difficult process than its unconscious homologous. However, there does not seem to be a reason to look at the child's theory as somehow unwarranted, as Fodor and Lepore would say:⁸² no more than any other theory is.

The other argument discussed by Fodor and Lepore seems to be connected with what we have already said:

If there is any fact of the matter at all about what the interpretation of a language is, then the evidence which selects a meaning theory for that language must in principle be *publicly accessible* data. [...]

The problem with this line of argument is that, though it may show that the evidence that determines the choice of a T-theory

⁸¹Cf. *ibid.*, pg. 225 n. 14.

⁸²*Ibid.*, pg. 81.

can't be "hidden" *in principle*, it doesn't begin to show that it can't be hidden from the child/linguist/radical interpreter.⁸³

But I think that this can be shown if we let the two lines of argument converge. For Fodor and Lepore seem to be, towards the requirements of the RI thought experiment, both too restrictive (when they claim that RI requires knowledge of just contingent behavioural facts) and, paradoxically enough, too liberal (when they claim that radical interpretation is possible for an omniscient being that knows all the nonintentional facts and what laws and counterfactuals are supported by those facts).⁸⁴ Truth, perhaps, lie in between: we need not know all true counterfactuals supported by facts, and even not know all facts. We need to have, though, some methodological principle which allows us to infer counterfactuals and laws from facts. Now, we know that something of the like is probably metaphysical in nature, and that, e.g., induction is never supported by facts. Indeed, this is enough to believe that even an omniscient being knowing all counterfactuals supported by facts would have more than one theory to choose from.

The question we should ask, then, is not whether radical interpretation is possible, but how to define radical interpretation in a reasonable way; that is, we have to specify better than Fodor and Lepore what knowledge the radical interpreter is allowed to have. This is my proposal: together with knowledge of contingent behavioural facts, radical interpreters are allowed to know facts concerning truth and meaning of their own languages. One might object that everybody has such knowledge. Indeed, radical interpretation is just an extreme example showing that the mechanisms explaining it are exactly the same mechanisms involved in our interpretation of our neighbors' discourses. And this is exactly what makes it matter to philosophy of language.

4.2.6 Nomologicity

We can now discuss the second of the three Davidsonian suggestions to solve the extensionality problem. This consists in considering T-sentences as laws – natural laws, in a sense. This should help solving the problem, because T*-sentences are not lawlike, since they cannot support appropriate counterfactuals to solve our twofold extensionality problem, i.e., distinguishing both between “‘Snow is white’ is true in English iff snow is white” and “‘Snow is white’ is true in English iff grass is green,” and between “‘ x is F ’ is true in \mathcal{L} iff x is F ” and “‘ x is F ’ is true in \mathcal{L} iff x is G ,” where F and G are two coextensive but not synonymous atomic predicates of \mathcal{L} .⁸⁵

Fodor and Lepore claim that T-sentences fail the test of counterfactuals as well, and therefore cannot be considered lawlike. On one hand, their being

⁸³*Ibid.*, pg. 80, emphasis in text.

⁸⁴*Ibid.*, pgs. 80 f.

⁸⁵Davidson (2001), pg. 26 n. 11.

laws would contradict the arbitrariness of language.⁸⁶ On the other hand, and for the same reason, the counterfactual world in which, say, “Particles of like charge do not repel each other” is true is one in which the conventions of English are different from here, and not a world where the laws of physics are different.⁸⁷

However, both the problems can be solved if we relativize T-sentence to languages, speakers and worlds. In this way our T-sentence will be “‘Snow is white’ is true at world w_1 in John $_{w_0}$ ’s language iff snow is white at w_1 ,” where w_0 is John’s world (probably the actual world), John $_{w_i}$ is the counterpart of John at world w_i (whatever our theory of counterparts is), whose language we are considering, and w_1 is the world were we are evaluating truth.

This solution, which seems more or less equivalent to the proposed amendment of Fodor and Lepore,⁸⁸ takes into account what they call the principle of conventionality. But here is the problem, according to Fodor and Lepore: radical interpreters cannot derive as laws our (amended) T-sentences, because they do not know about the principle of conventionality.⁸⁹

I think that this conclusion holds only in one formulation of such principle. If we define such a principle as stating that languages are essentially arbitrary in their relation to the world,⁹⁰ then we should doubt that such an assumption can be made by the radical interpreter. If, on the other hand, we think of the principle of conventionality as stating that each speaker has an idiolect for which a peculiar theory of meaning holds, then we do not see why we should not assume knowledge of this by the radical interpreter. For it merely translates in this requirement: assume that the radical interpreter does not know from the beginning the theory of meaning for \mathcal{L} , and that from the beginning there is no reason to expect such a theory of meaning to be the same as one that the radical interpreter already knows. It seems to me that this is a fair restatement of the radical interpretation requirement, and that, therefore, it is fair for us to ask that it be satisfied by the radical interpreter, whose task is defined by that requirement.

There is one objection by Fodor and Lepore that still seems to hold: the case of children. It seems fair to say that while field linguists can very well accept the prudential principle of not assuming that all people they come in contact with are speakers of the same language, children do not know about such a principle. In fact, they acquire a unique language, while were they to assume that every other speaker has a unique language, they should acquire the (overlapping) grammars of all the speakers they come in contact with.

Of course, the two situations would have more or less the same cognitive

⁸⁶Fodor and Lepore (1992), pgs. 84 f.

⁸⁷*Ibid.*, pg. 85.

⁸⁸*Ibid.*, pgs. 86 ff.

⁸⁹*Ibid.*, pg. 89.

⁹⁰Which, by the way, should make us dub this “principle of arbitrariness,” since conventionality is another thing. Cf. Davidson (2001), chpt. 18.

output, but let us assume, for the sake of argument and because it seems more plausible, that the child from the beginning has one language, and according to the theory of meaning of that unique language interprets what is said by all other speakers. In Fodor and Lepore's words:

It's only when they are told about there being lots of languages other than English that they understand that there are nearby worlds in which "Snow is white" is true even though snow isn't white.⁹¹

To show that there is no contradiction in this case, we need to introduce the principle of charity. I still postpone its discussion to § 4.2.7, but here is a sketch of what I have in mind: suppose that the radical interpreter was not to assume from the beginning that \mathcal{L} is an unknown language. Suppose, in fact, that interpretation begins by trying to match \mathcal{L} with one's own language, and only when this proves impossible because of some empirical evidence, the interpreter tries to adjust the theory of meaning so to match the data. In the case of children, this is our explanation: children assume that everybody speaks the same language, and acquire that, as if it was one language, unless too much contrasting evidence leads them to prefer acquiring different sets of rules. When interpreting, then, they first try to match what they are told with the set(s) of rules they already know, and only after that they try to sort out the problem by making adjustments, as slight as possible.

The first problem that Fodor and Lepore note, even with such an elaborate adjustment to the nomologicity suggestion, is similar to the one noted above with regard to compositionality. Assume that "Water is H_2O " is a law. Then there is no way to distinguish between the meanings of "Water is wet" and " H_2O is wet."

This problem, too, is solved by considerations about charity. In fact, it is perfectly acceptable to assume ignorance of the technical term H_2O , while the term "water" is considerably more common. In translating a language that shows only one of the two terms, then, the latter translation is to be preferred. In translating a language with two different terms, the latter will be associated to the most common in \mathcal{L} .

The second consideration put forward by Fodor and Lepore is that it is not at all clear that nomologicity entails holism. An argument from the former to the latter should prove that the only way to give nomological strength to T-sentences (i.e., to amended T-sentences) is by taking into consideration different sentences of \mathcal{L} . If this cannot be proved, we have no argument for the impossibility of punctate languages and, therefore, for holism.⁹²

⁹¹Fodor and Lepore (1992), pg. 89.

⁹²*Ibid.*, pgs. 91 f.

I agree with the conclusion that nomologicity alone is not enough to prove holism. But I think, as it should be clear from the present section, that I also think that nomologicity needs, to be true, the principle of charity, and there are reasons to believe that the principle of charity does entail holism.

Nomologicity, anyway, is not pointless: it plays the final role of our semantic explanation. From the theory of interpretation grounded in methodological principles to a theory of meaning a step is necessary, a step to rid us of underdetermination of theories of meaning with respect to behavioural evidence. I think that this step is accomplished by a nomological principle, which enables us to treat the T-sentences found through a theory of interpretation as laws which support counterfactuals. Perhaps some other underdetermination will be at stake: the Q/D thesis probably holds in this case. But, if this is true, then theories of meaning are in the company of all empirical theories, which is exactly the company we want them in from the beginning.⁹³

4.2.7 Interpretation with Charity

We now turn to the discussion of the third Davidsonian suggestion: the principle of charity. Faced with the intrinsic underdetermination of semantic theories (i.e., of the underdetermination that semantic theories have inasmuch as they are empirical theories), Davidson proposed his own principle of charity to dispose of some hypotheses without verifying them through experience. He believed Quine to be too restrictive when he called evidence just the speaker's behaviour. There is, in fact, also our own linguistic behaviour, and according to it we can interpret others' behaviour as non-deviant.

Here is Davidson's proposal:

A good place to begin is with the attitude of holding a sentence true, of accepting it as true. This is, of course, a belief, but it is a single attitude applicable to all sentences, and so does not ask us to be able to make finely discriminated distinctions among beliefs. It is an attitude an interpreter may plausibly be taken to be able to identify before he can interpret, since he may know that a person intends to express a truth in uttering a sentence without having any idea *what* truth.⁹⁴

Holding true, thus, is the crucial attitude to begin interpretation. We can identify such an attitude, because we know what it is like to hold a sentence true, and because we are assuming that the speaker we want to interpret is, under this respect, akin to us. But to attribute someone this attitude we need to attribute them a whole lot of beliefs and attitudes similar to ours. In fact,

⁹³More on the subject of semantics as theory of meaning, as opposed to theory of interpretation, will be said in the next chapter.

⁹⁴Davidson (2001), pg. 135, emphasis in text.

there can be no definite list of things that must be believed by someone who understands a sentence and holds it true. Endless connected beliefs might be necessary.⁹⁵ Therefore it is clear that, unless our first move is to attribute such beliefs and attitudes, we cannot make it intelligible that our subjects hold something true.

But the point is that, according to Davidson, the principle of charity is not only useful for our interpretation of other people's speech: it is also necessary. If we do not assume that the cognitive system of the speakers we are surveying has almost the same characteristics as ours, we cannot even consider it a cognitive system, nor consider them as speakers, as opposed to incomprehensible babblers. Therefore, we would lack reasons to decide to construe their utterances as carrying some meaning. The principle of charity, hence, is not mere prudence: it is essential to our mastering of language as a means of communication.

We must obviously apply this criterion *cum grano salis*, for the speaker can after all be mistaken.⁹⁶ Anyway, mistakes can be intelligible only if most beliefs are actually shared. We need a common background to indicate what the different beliefs are.⁹⁷

We have seen that both within a behavioristic theory or within a theory grounded on a strong principle of charity, the very first move to do is to determine whether the speakers hold their sentences true or false. In Quine's theory, this leads to a trial and error procedure to determine possible meanings of the sentences under scrutiny, taking other sentences and other such procedures into consideration. In Davidson's theory, the attribution of such an attitude leads to the determination of meaning, because even an attitude like that of holding a sentence true cannot be intelligible unless we assume the common human background of beliefs and attitudes, of which holding true is a part.

Belief Attribution

One way of understanding how assuming the principle of charity entail semantic holism is to see it as entailing an inference from belief holism to semantic holism. The system of beliefs appears to be holistic because belief interpretation is constrained by the principle of charity, which entails holism.⁹⁸ In fact, in order to have a certain belief (or other propositional attitude) it is necessary to have other propositional attitudes, for we can interpret someone as having a belief only if we can interpret that belief as being related in a

⁹⁵Cf. *ibid.*, pg. 158.

⁹⁶*Ibid.*, pg. 136.

⁹⁷*Ibid.*, pgs. 199 f.

⁹⁸Fodor and Lepore (1992), pg. 113.

certain way to other intentional states.⁹⁹

The movement from belief holism to semantic holism is made by what Fodor and Lepore call primacy of belief thesis:

The conditions for content attribution *inherit* the conditions for belief attribution; hence, if the former are holistic, then the latter must be too.¹⁰⁰

Note that the name “primacy of belief” is misleading, because, as it will become clearer later, what is claimed is that the individuation of intentional states comes before the attribution of content. This, if true at all, is true for beliefs as well as for desires and speech acts. With this caveat in mind, we can stick to Fodor and Lepore’s phrase.

The reason for believing in the primacy thesis, *prima facie*, is that it seems, thanks to the principle of charity, that the attribution of beliefs is essential to attribute semantic content to utterances. In other words: only if the speakers have certain beliefs, their utterances have certain meanings. This is because uttering something depends on performing certain speech acts, which depends on its turn on having certain propositional attitudes.¹⁰¹ And, as we have said, only in the context of a system we can identify propositional attitudes, because propositional attitudes are functionally identified by the relation they bear to other mental states and, specifically, to other propositional attitudes.¹⁰²

Fodor and Lepore are willing to concede that much, but they claim that neither the primacy of belief thesis follows, nor, therefore, semantic holism. In fact, semantic holism would follow only if the functional role that is essential for the identification of beliefs (and, in general, of intentional states) is also what determines the intentional content of beliefs. Fodor and Lepore propose an alternative: content derives from mental states, *viz.*, mental representations.¹⁰³ If the identification of mental representations is not a functional process, then holism cannot be inferred:

So, then, to summarize: the species of mental representation theory we have in mind [...] tells one kind of story about the difference between believing and, as it might be, wanting and quite a different kind of story about the difference between believing that P and, as it might be, believing that Q. The former story is *bona fide* functionalist: believing and wanting are both modes of entertaining mental representations, but they differ with respect to the causal (associative/computational) roles that the mental

⁹⁹ *Ibid.*, pg. 107.

¹⁰⁰ *Ibid.*, pg. 114, emphasis in text.

¹⁰¹ *Ibid.*, pg. 122.

¹⁰² *Ibid.*, pgs. 115 ff.

¹⁰³ *Ibid.*, pgs. 122 f.

representations play (that is, with respect to how the representations interact with inputs, outputs, and one another). By contrast, the difference between believing that P and believing that Q invokes the semantic properties of mental representations, and these are supposed to be grounded in mind/world relations *rather than* functional roles.¹⁰⁴

I would rather say that the difference is one of mind/world relations, *that is* of functional roles. Fodor and Lepore would want their mental representations to have a direct causal relation with “the things they denote and the properties they express.”¹⁰⁵ This is the reason why they judge that content attribution is not a functional process:

Whatever bestows semantic evaluability upon mental representations, it can’t be that they are used to express beliefs or intentions. Mental representations aren’t “used” at all; and we typically have no beliefs or intentions with respect to them.¹⁰⁶

If this is true, I wonder what mental representations are doing in the theory at all. Whenever we investigate our thoughts, we never find such mental representations as distinct from beliefs or intentions. On the other hand, consider how the property of being a determined mental representation of *P* gives rise to a belief about *P* or an intention about *P*. This difference must arise from the different functional role played by the mental representation. But if this is so, we can assume that the same functionalistic mechanism is responsible for both the content and the form of the intentional state. Another passage by Fodor and Lepore is worth quoting, in order to get my point sharper:

Maybe thinking that the cat is on the mat is just saying to oneself “The cat is on the mat.” However, if it is, then “saying to oneself” must differ from “saying” *tout court*. For saying *tout court* is normally the consequence of deciding what to say; that is, of decision-theoretic processes which invoke beliefs and desires. Whereas, saying to oneself can’t presuppose believing or desiring, since, according to the present account, believing and desiring are themselves species of saying to oneself.¹⁰⁷

I would like to reply that, “normally,” “saying *tout court*” is anything but “the consequence of deciding what to say.” In most cases, in fact, conversation

¹⁰⁴ *Ibid.*, pg. 127, emphasis in text.

¹⁰⁵ *Ibid.*, pg. 125.

¹⁰⁶ *Ibid.*, pgs. 125 f.

¹⁰⁷ *Ibid.*, pgs. 126 f.

is carried out almost in an automatic way, especially in the case of largely ritualized conversations, like greetings, ceremonies, narratives, and so on. These acts surely are very important when we want to specify possible uses of our language.

Conversely, we can ask a reason for each of these acts. And a reason is an explication in terms of beliefs and desires. It is the same with saying to oneself. There seems to be no reason for a thought to come in our mind instead of another, but are we sure that scrutiny would fail to reveal any? Usually, an accurate inquiry leads always to reasons for any thought: the causal origin of any mental representations is grounded on other mental representations. Indeed, if the causal relation between mental representations and denoted objects were as Fodor and Lepore imagine it, it would be determined by the object which mental representation is aroused. In fact, it is not so: different people have different mental representations towards the same objects.

It can be said that this happens because of different physical conformations that alter the causal connection. But then we should explain why different people come to have the same mental representations toward some objects and different ones toward some others. It seems that some intentional explanation must enter the picture, somehow.

Every linguistic act (like every act) has physical causes, but if we examine the relation between the act and its cause, we realize that, to characterize the act as an act, there must be a rational component between them. If we now consider the origin of the rational component, another intentional element must reveal itself, together with a physical cause.

The problem is difficult because here tradition considered to be dealing with two distinct and opposite areas, while it is probably a continuum. To be more precise: before any theoretic determination, there is no principled difference between an act and an event. After theoretic determination, there is a border, which can always be moved further by further investigations; but there is an end to investigations. When we reach it, we have two options. Either we try to root all the chain to the external world by way of causal relations, or we ground it in human rationality. The first route is difficult to pursue, because the end we have reached is arbitrary, so there is no possible justification for causality to attach exactly to that point of the chain and not to another. The other route, though, presupposes an analysis of human rationality and human nature, which I will not attempt here.

Suffice it to say that rationality, under this perspective, works like a regulative principle: unless we consider the human intentional system as complying with our ideal of rationality (which is holistic, as it states that it is possible to find reasons for every intentional state), we fail to attribute intentional content to human mental states.¹⁰⁸ In fact, nothing warrants us that human actions are anything more than a mere causal response to external stimuli. However,

¹⁰⁸More on this subject will be added in § 4.3.3.

this hypothesis is incompatible with our actual interpretative practices.

To this line of argument it might be objected that we could fail to appreciate certain differences (e.g., between a belief and another) that nevertheless exist:

The problem is that if the Projectivist account of the (putative) interpretive element in belief attribution is right, then what *you* can believe depends on what *your interpreter* can say. But if anything is metaphysically independent of anything, surely your repertoire of potential beliefs is independent of anybody else's repertoire of potential speech acts. There is, no doubt, an "element of interpretation" in talk about mountains; where does the mountain end and the valley begin, after all? But only a megalomaniac could suppose that whether there are mountains depends on whether he can say that there are.¹⁰⁹

Indeed, my potential beliefs do not depend on my interpreters' potential linguistic expressions, but the intelligibility of my beliefs to them – i.e., what they can think about me and my belief – does. In a sense, it exists (if it exists at all) only what we can say it exists, because we cannot even wonder about the existence of something else. The world's ontology does not depend on vocabulary, but the part of it we can know and understand does.

Fodor and Lepore try to build a counterexample:

Consider the following kind of case: Smith's three-year-old hears him prattling about the analytic/synthetic distinction, and it occurs to the child that Smith must have some beliefs about this distinction that he, the child, does not understand and could not express. On the present analysis, this thought – which intuition might plausibly take to be true – is self-contradictory, since it entails both that Smith is in some state that would normally lead the three-year-old to say that "blah, blah, blah, analytic/synthetic, blah, blah, blah" and that there is no such state.¹¹⁰

But the point here is not that there is no state which would lead the child to utter that, but that the child does not know which state it is, and therefore cannot correctly interpret Smith's words. On the other hand, because some (extended) principle of charity is applied, the child still knows that there is such a mental state, and that Smith's words are, in fact, meaningful, and not mere prattling. Indeed, even the latter option is open to the child, and it often happens that children, and sometimes adult people, take for insensate prattling what they do not understand. And sometimes to assume that a

¹⁰⁹ *Ibid.*, pg. 140, emphasis in text. Projectivism is the label that Fodor and Lepore attach to a theory of interpretation along the present lines.

¹¹⁰ *Ibid.*, pg. 139.

discourse is meaningless is just the right move, because sometimes people do consciously utter nonsense – e.g., for amusement.

Another objection advanced by Fodor and Lepore is the following:

If there's a law that makes being in intentional state A nomologically sufficient for being in intentional (and/or) behavioral state B, then, given the knowledge that a creature *is* in state A, you can predict that it will (come to) be in state B, *whether or not the transition from A to B is rational*. The upshot is that the argument that infers charity about rationality from the presuppositions for intentional prediction fails because it *begs the question against there being intentional laws*.¹¹¹

But intentional laws are recoverable only if some principle of charity is assumed. In fact, we can connect two intentional states *A* and *B* by a law even not knowing how the transition from one state to the other goes, but it is only because we believe that, were we to inquire on this issue, we would find some reason, i.e., some rational connection,¹¹² that we take this connection as lawlike and supporting counterfactuals and previsions. None of the examples put forward by Fodor and Lepore to support their thesis points in a different direction.

The last objection against the possibility of inferring semantic holism from the practice of radical interpretation, has a skeptical flavour:

But, one might suppose, it's always possible that the interpreter should have *misidentified* the causes on which the informant's holding-trues are contingent – in which case, the interpreter might operate on the principle of making the informant's judgments true by the interpreter's lights, and yet, for all we know, the informant might be saying something that is actually false. It looks as if the principle of charity leaves open the possibility of *folie à deux* (which is, of course, just what the skeptic always thought).¹¹³

I have no metaphysical problem with the acceptance of skepticism on this issue, as I have already made clear. But consider two things: first, future interactions between the interpreter and the informant will reveal misunderstandings and misinterpretations. Second, even if the principle of truth (i.e., the principle according to which the speaker utters mostly truths) is untenable, still there is a role played by the notion of truth that is useful for us. This is not the idea of truth as the absolute objectivity to which we must compare

¹¹¹ *Ibid.*, pg. 152, emphasis in text.

¹¹² Cf. Davidson (1980), pg. 233.

¹¹³ Fodor and Lepore (1992), pg. 159, emphasis in text.

what we say, but as a kind of intersubjective space where we, as human beings, can interact. For everything that matters to us, this is the only kind of truth we should care: what is held for true by human beings. This, together with the principle of charity, gives us enough room to be wrong about the world, to find our mistakes, and to correct them. Which is fine, because we happen to make some mistakes from time to time – that is, we sometimes happen to undermine interactions.

It should be clear, now, that the notion of truth is central for these holistic theories, while those of reference and meaning can only be derivative. To attribute meaning or reference to a sentence or its constituents before having decided (albeit provisionally) its putative truth value would oblige to make way more assumptions about the contingent status of the speaker's cognitive system. In other words, we could not limit ourselves to determining whether the speakers hold their sentences true or false, but we would have to determine, as a start, what they are talking about and what they are saying about it. Such determination would be a very difficult task to accomplish, and would increase dramatically the risk of error.

For a Correct Understanding of the PoC

Some of the problems that Fodor and Lepore have with the principle of charity depend on its definition as they construe it:

[...] truth conditions must be assigned to formulas of L under the constraint that most of the sentences held true by a speaker of L *are* true. This principle is supposed to imply content holism on the intended interpretation, which is that “most of the sentences” means a lot of them.¹¹⁴

This is a slightly incorrect way of putting things. Here is how Davidson stated the principle:

This method is intended to solve the problem of the interdependence of belief and meaning by holding belief constant as far as possible while solving for meaning. This is accomplished by assigning truth conditions to alien sentences that make native speakers right when plausibly possible, *according*, of course, *to our own view of what is right*. What justifies the procedure is the fact that disagreement and agreement alike are intelligible only against a background of massive agreement. [...]

The methodological advice to interpret in a way that optimizes agreement should not be conceived as resting on a charitable assumption about human intelligence that might turn out to be false.

¹¹⁴ *Ibid.*, pg. 93, emphasis in text.

If we cannot find a way to interpret the utterances and other behaviour of a creature as revealing a set of beliefs largely consistent and true *by our own standards*, we have no reason to count that creature as rational, as having beliefs, or as saying anything.¹¹⁵

Under this interpretation, holism is of course entailed. And it is quite evident that a principle of compositionality is entailed too. This is because to have the background of massive agreement required by Davidson one must obviously consider more than one single sentence, and perhaps a lot of sentences, which is enough to satisfy Fodor and Lepore's definition of holism. The principle of compositionality, which provides the systematicity which we require under our definition of holism, is entailed because we cannot assume agreement concerning a term, unless that term is consistently used in agreement with our own usage.¹¹⁶

Thus, of the two points against the utility of the principle of charity (PoC, as Fodor and Lepore abbreviate it), the first (the fact that it would not prevent punctate languages) is already defused. The second claim is that the PoC is in fact pointless, since it is the merely repetition of "the truism that a good theory had better comport with the data."¹¹⁷

In fact, it is more than that. The PoC allows us to give a first interpretation of the data, because it tells us that if what we are faced with are utterances, then we better interpret them in agreement with our beliefs. When we want to interpret a speaker of \mathcal{L} , first of all we look for the utterances held true by speakers of \mathcal{L} . Then, we assume that they are right in holding true their utterances, and look for the sentences that we would have held true in the same circumstances, and identify the meanings of the sentences. Of course, this process is subject to error.

The PoC intervenes several times in this process: first of all, in choosing to identify an attitude which we can recognize, we are implicitly assuming the rationality of alien speakers. Second, in tracing a correspondence between our linguistic behaviour and those of the alien speakers, we are employing the principle of charity. Finally, in assuming that we can sort out the errors, and make the alien speakers rather mistaken than irrational, we are holding onto the PoC. Of course, if one assumes the PoC to be formulated in Fodor and Lepore's way, none of these uses is allowed, and the principle indeed reduces to a truism.

Fodor and Lepore carry on, attempting to prove that the PoC is not even sufficient to solve the extensionality problem, which is the reason it was introduced in the first place. They build an elaborated example to explain their

¹¹⁵Davidson (2001), pg. 137, emphasis added. Cf. Quine (1960), pg. 69.

¹¹⁶Cf. Davidson (2001), pg. 168.

¹¹⁷Fodor and Lepore (1992), pg. 96.

point.¹¹⁸ The problem with it, without going into details, is that they are not explicit in the difference between object language and metalanguage, so that it is never clear if the pair of nomologically coextensive but not synonymous terms “white” and “*F*” that they mention belongs to the former, to the latter, or both.¹¹⁹

However, the problem with their appeal to pairs of nomologically coextensive but not synonymous terms is general: how can the speaker be sure that the difference between *F* and white is clear to any interlocutor, if in no context and in no verbal explanation such a difference can be explicit? The speaker is always at risk of being misinterpreted. Moreover, the Wittgensteinian argument against private language applies, because the speaker can never be assured of being consistent with a difference that is never explicitable. In such a situation, the two terms come to be synonymous, because this is the way they are used, and if the theory of meaning has to comport with the data, this is precisely what the data show. I will repeat myself: nothing forbids to assume that words and terms have a meaning that is never shown by their use, except economicity of explanation.

The last possibility that Fodor and Lepore contemplate, “in a spirit of comprehensiveness,”¹²⁰ is the following:

Suppose there are expressions which (1) don’t appear in token-reflexive utterances (or don’t appear in them often enough to matter) and (2) which are *syntactically* atomic. Well, by assumption, the interpretation of these expressions can’t be fixed by their behavior in token-reflexive utterances; and since they are syntactically atomic by assumption, their interpretations can’t be determined compositionally from the interpretations of their parts. We don’t know whether there are such expressions, but we don’t see why there couldn’t be. Very theoretical terms like “proton” would be among the likely candidates. The suggestion is that it is in the radical interpretation of *these* sorts of expressions that appeals to charity somehow play an ineliminable role.¹²¹

However, Fodor and Lepore claim, if the interpretation of such terms is achieved thanks to their connection to other terms which do appear in token-reflexive utterances, no charity is needed, because interpretation of token-reflexive utterances does not need charity. And if no such connection is given,

¹¹⁸ *Ibid.*, pgs. 98 ff.

¹¹⁹ Also the fact that they are never able to furnish a single example from a real natural language seems telling to me. Although it does not diminish the logical importance of their argument, still it gives the impression that in actual interpretation with real languages this sort of problems are in fact solved.

¹²⁰ *Ibid.*, pg. 100.

¹²¹ *Ibid.*, pgs. 100 f., emphasis in text.

if therefore the term appears only in purely theoretical contexts, then it is precisely in such contexts that we expect some disagreement and, therefore, the principle of charity not to apply.¹²²

This is certainly true, but again it is just a consequence of having misunderstood the principle of charity. If we take it back to the methodological status that is its own, we can see that it is precisely because of the principle of charity, and not in spite of it, that we assume highly theoretical contexts to trigger disagreement. And it is because of the principle of charity and of the empirical enterprise that has it as a methodological principle that we start looking for token-reflexive utterances in the first place, and that we know how to identify them (they are, usually, those whose truth conditions are related to the context of utterance).

4.3 Interpreting One's Own Sentences

We have seen how, according to the account given in the last section, semantic holism is able to give us empirical hypotheses to interpret others' sentences. Interpreting is a procedure that starts with attributing our own system of beliefs to interpret other's sentences as if they were produced by ourselves, and then proceeds making the necessary changes to reduce as much as possible the attribution of incoherences and irrationalities. The behavioristic account, though more indeterministic in its results than the other one, is not dissimilar from it: the mention of a system of beliefs can indeed be avoided, though perhaps not completely replaced, through behavioristic means. Whether we stick to one or the other theory, though, we need some account of our intuitive comprehension of our own sentences.

4.3.1 An Atomistic Knowledge?

This is the place where an atomistic account may seem to intervene: although we need holistic means to interpret others' speech, we may think that the meaning of every sentence we produce is known to ourselves through one of the ways we have seen in § 2.1. After all, the speaker is the one who produced those sentences, and therefore knows their meaning. But this move sounds illegitimate:

Perhaps someone [...] will be tempted to say, 'But at least the speaker knows what he is referring to.' One should stand firm against this thought. The semantic features of language are public features. What no one can, in the nature of the case, figure out from the totality of the relevant evidence cannot be part of meaning. And since every speaker must, in some dim sense at least,

¹²²Cf. Kuhn's remark about incommensurability, quoted above in § 4.2.2.

know this, he cannot even intend to use his words with a unique reference, for he knows that there is no way for his words to convey this reference to another.¹²³

In other words, since language has an effective use as a means of communication, and since the exactness of reference individuation cannot ever be reached, we must conclude that the latter cannot have any importance for the success of communication. By saying this, I mean that an element of indetermination always remains, but this indetermination only occasionally hinders communication. Hence, there is no benefit for the theory in assuming that meaning is known to the speaker in an atomistic way.

Let us try, then, to explain with other means how speakers know the meanings of their own sentences. Odd as it may sound, I believe some kind of principle of charity is at work in this case, too.

4.3.2 Social Speakers

One characteristic of language that we have neglected so far is that language is social. I do not mean that it is not possible to use language in private, but I think that it is better to study it when it is exploited for social purposes.¹²⁴

This is the reason behind the principle of charity. Since language is a social instrument, and a society can be established only among peers (i.e., people judging each other on the same level with respect to some relevant quality or characteristic),¹²⁵ we must attribute those qualities or characteristics we deeply feel as ours to those we wish to understand. In fact, language is a characteristic universally attributed to human beings. Although there might be human beings completely cut out of the linguistic interaction,¹²⁶ it is not them we think of first when we think about human beings: we think of people we can speak to.

Not only, in relating to others, we put them on the same level as ourselves: conversely, we also attribute to ourselves the same characteristics that we attribute them. In particular, this is true for what concerns one's linguistic conduct. Just as we reconstruct the meaning of other people's speech by assuming that their sentences are coherent with their beliefs and motives, we can look at our sentences as being coherent with our beliefs and motives, and

¹²³Davidson (2001), pg. 235.

¹²⁴This is not equivalent to claiming that the purpose of language is communication, which I think it is not (cf. Chomsky (1980), pg. 230). After all, communication is only one form that social interaction may assume. Anyway, I will not pursue this issue any further here.

¹²⁵I keep vague, here, because societies can come in great variety.

¹²⁶I stress that this is strongly dubitative: we talk to human non-speakers in a variety of occasions, and we often assume that there is some interaction mediated by language, even though we are quite sure that there is no comprehension in the narrow sense of the word. Some people talk to relatives in a coma, and most cultures contemplate speaking with the dead. Although the declared motivation is sometimes to attain self-comfort, there is often the (sometimes not so) hidden motivation of being sensed, if not heard and listened at.

thus reconstruct their meaning. And, again similarly to the case of interpretation of others' speech, we can correct our mistakes, when we realize that our words had indeed a different meaning from that which we wanted them to convey.

In other words, we interpret other people's sentences because we understand our own language, and we attribute it to them, or at least we attribute to them the system of beliefs and motives that produces it. We understand our own language, then, because we attribute to ourselves that very system of beliefs and motives. Our next question, therefore, must be how this latter attribution is possible. I think that, again, we should look first at how this is possible towards other people, and then try to apply what we have found to ourselves.

It is true that we start with assuming other people's systems of beliefs and motives. But this is possible, first of all, because *they*, at least sometimes, *do tell us about their beliefs and motives*. I want to suggest that people know that their beliefs and their motives are there, and what they are, because they can verbalize them. We take notice of the content of our mental representations – and of their existence in general – only when they take part in the verbalized system of our linguistic acts.

To assume the existence of mental representations even when we cannot verbalize their content is totally irrelevant for the present argument. In fact, even if unconscious representations exist, they can have an effect on our conscious life either inasmuch as non-representational causes, or inasmuch as representations. In both cases, direct reasoning and consciousness are not supposed to interact with them, for they are unaware of their existence. Thus, we cannot determine from the effect they bear on our conscious life what their nature and the nature of their interaction with our conscious mind really are. In other words, we cannot infer from the presence of some unconscious effect the existence of unconscious mental representations. If we want to assume such existence (for economy of explanation, maybe), of course we can: we are allowed to do it by the existence of conscious mental representations, which we are assured of because of conscious awareness and interaction, of which verbalization is a privileged form.

4.3.3 A Rational System

One could object that the explanation I give is, in fact, a vicious circle, since it explains linguistic meaning resorting to beliefs and motives, and beliefs and motives resorting to linguistic meaning. So we should state clearly what comes first.

I think that the right answer is that neither comes first. The two elements at stake – i.e., the semantic system and the system of beliefs and motives – are not two separate stages of a procedure that goes necessarily from one to the other and cannot go back. They are parts of a whole system, which we

can dub ‘rationality,’ or ‘rational system.’

Each one of us, somehow, is firmly convinced that the ancient definition of man as a rational animal is, more or less, right. This definition may seem not quite right concerning other people, but everyone knows that it is right for herself or himself. I do not mean that people are always rational, of course, nor that everyone judges himself or herself as having been always rational. I mean that, looking retrospectively to events in our lives, we can rationally reconstruct our conduct so that we can regard it as the product of a rational being.

In other words, our own sentences have a meaning not because they are formulated in some abstract language whose semantic rules we happen to know, but because that language is our language; we know their meaning because we have produced them, and this production is part of our life. In this sense, the real primitive of such a holistic theory of interpretation is rationality. Every human being is rational (and this explains how each one can produce meaningful sentences in the first place), and can attribute such rationality to others (and this explains how each one interprets the meaning of the sentences others produce). Rationality, then, is axiomatic: no one can question it, at least for what concerns herself, and those she chooses to speak with.

CHAPTER 5

Holistic Meaning

Interpretation is more a psychological or anthropological matter, rather than a properly semantic and linguistical one. Despite a tendency in modern linguistics to consider the faculty of language specific of human beings, it is very rare that the study of linguistic mechanisms and laws themselves makes any appeal to this specificity. The answer to the properly semantic question I mentioned above can vary, depending on what opinion one entertains about the relation between language and human nature. Therefore, I will now discuss this issue.

5.1 The Place for Linguistics

There are scholars who programmatically choose to deem language a kind of structure which could be implemented by different cognitive systems rather than human:

The main question I will attack in this paper is the one that is really central [...]: it is the question as to whether the role played by Universal Grammar can be assigned to a completely abstract and, we might say, platonic object, one that is unrelated to any particular biological characteristics of human beings.¹

Nothing in the nature of the human being, as biology and evolutionary theories describe it, implies that the faculty of language is necessarily part of it. In fact, according to those theories and disciplines, there is nothing necessary about human nature and its history.

Conversely, according to Goldsmith, nothing in the nature of language as linguistics describes it implies that it is necessarily a human being that possesses and exploits it. He claims that his position is not at all different from the original generative one.² Anyway, there is certainly a physicalistic

¹Goldsmith (2007), pg. 13.

²*Ibid.*, pgs. 13 and 31.

turn in Chomsky's views and in those of his followers, which is opposed by Goldsmith:

[...] like every other linguist, Chomsky and his colleagues are not capable of establishing where and how their theories of grammar are instantiated in the brain; but they interpret their failure to do so not as a challenge to physicalism, but as an IOU, that is, a promise that at some unspecified future date, a physical location in the brain will be found. *There is no difference between promising, some day, to do something in a yet undiscovered way and not promising to do it at all.* It is not the linguist's job to determine how the brain works: that is a good thing, since there are few linguists with any serious training in neuroanatomy. It is the linguist's job to figure out how language works, and as that challenge continues to be handled, linguists and neuroscientists will be able in the future to come up with a synthetic view.³

Goldsmith may perhaps be misled in one point. After all, very few linguists (if any at all) have their agenda written by neuroscience, and Chomsky is not among them. Besides, there is no reason to wait until a further development of linguistics to let it interact with other scientific enterprises. All sciences, in fact, are developing and changing all time, and there is no chance to foresee the complete development for even one of them.

Even if Goldsmith may be wrong in his attack against Chomsky and the physicalist turn he imposed to generative linguistics, he gets one point: although the study of language should not isolate itself from other disciplines, it needs not depend on them, either. Perfectly legitimate theses in linguistics and philosophy of language can deal with language as a structure, independently on the way such a structure is instantiated, and to what extent.

Such a position about autonomy of linguistics could be interpreted in two different ways. In one reading (call it the weak reading), linguistics is autonomous with respect to other disciplines in the sense that it is not submitted to restrictions from those disciplines, but it cannot be excluded that linguistics itself may set restrictions to, for instance, neurological or psychological researches. In another reading (call it the strong reading) the object of linguistics, i.e., language, is radically different from the object of other disciplines, and no interaction is possible.⁴

It is remarkable that, given one linguistic theory, it can usually be read and interpreted according to both readings. This is especially true for what con-

³*Ibid.*, pg. 30, emphasis in text.

⁴In April 2007, Goldsmith gave a lecture in Pisa, on the subject "Towards a new empiricism in linguistics." When, in that occasion, I asked him about this double reading allowed by his views, he told me that he could not say which interpretation better described his position.

cerns syntax. There are indications that the same thing holds for phonology or morphology.⁵

Although this can be generally true also in the case of semantics, some specification is needed. Leaving the details for later investigation, it suffices for the moment to say that we may think the semantics of a particular language either to be an autonomous system, or to have strong relations with some conceptual system or conceptual structure, which is yet to specify.

According to the first alternative, to understand the mechanisms and laws governing the semantics of our language, we need not make any assumptions concerning human psychology or neurology. This could be a very much welcome idea, since there is little agreement about the former, and the major agreement about the latter is that we are probably only beginning to understand how it functions. The second alternative, therefore, is relatively unfavoured by semanticists: in spite of their programmatic declarations, they tend to resort very rarely to disciplines external to their own for a solution to key problems. In the rest of this section I will try to describe what holistic semantic theories of both kinds might look like.

5.2 Formal Semantic Theories

Let us examine first an instance of a semantic theory that starts from a perspective which traces a boundary between the linguistic and the conceptual features of meaning. It is fairly intuitive that whatever we say means what it means for two different reasons: one is the language it is formulated in, and the other is what we know about the world – i.e., at least that part of the world we are talking about. We want now to examine our question about semantics under the hypothesis according to which the two systems implied by these two orders of reasons (i.e., the linguistic system and the conceptual system) do not work in strong connection, but are, by and large, independent.

5.2.1 The Two Components of Meaning

First of all, we should ask ourselves what makes such a hypothesis sound reasonable and plausible. Consider the sentences in (1):⁶

- (1) a. This is too little carpet for the money.
- b. There are three wines in the cellar.
- c. Cat came.
- d. The three Kims I met yesterday were all tall.

⁵Cf. Celata and Calderone (2008).

⁶For the examples and a discussion in part along these lines, see Borer (2005a), pgs. 8 ff.

These sentences show some very well known facts about English, viz., the possibility of exploiting count nouns as mass nouns, and viceversa, and the possibility of exploiting proper nouns as common nouns, and viceversa. It is easy to see that **coercion** – i.e., the overriding of the semantic properties that we usually associate with terms such as ‘carpet’ or ‘wine’ (in this case, their being count or mass respectively) – is carried out by the particular syntactic context in which such terms are used.

We may make the hypothesis that the meanings of the terms are given by the specification for each term of the syntactic contexts in which it can occur, together with the value that it will assume in each of those contexts. Such a hypothesis, though, would not be very economic. In fact, there is a huge quantity of contexts in which an individual term might occur – perhaps an infinite number (depending on what we consider to be the relevant syntactic context).

Since we are imputing to syntactic contexts the shifts of meaning here under examination, there is no need to attach the knowledge of syntactic contexts to the knowledge of terms: we can just accept that syntax itself is responsible for meaning. In other words, we need not specify for each term the value it assumes for every syntactic context, because that value depends entirely on the syntactic context; therefore, syntactic rules can account for it.

It might be objected that, of course, there is something in the meaning of one term that does not depend on its syntactical function. This is certainly true, but we have to be more detailed than that in specifying the extent of the non-syntactic component contribution. Consider (2), for instance, and what Borer has to say about it:

(2) The red under fived lunch.

While Harley and Noyer actually cite (2) as an impossibility [...], it turns out that native speakers of English are perfectly capable of assigning an interpretation to it (some creature with some ‘bottom’-related properties, for example, a bottom-dweller, which is red, ate lunch five times; multiplied its lunch by five; divided its lunch by five; etc.). Of considerably more significance is the fact that (2) is not a word salad. It has an understood actor, *the red under*, who acted in some manner pertaining to *five* on some target, *lunch*.⁷

What this example shows, then, is that even the kind of ontological entity associated with a term depends on its syntactical position. In the syntactic context of (2), in fact, ‘five’ is an action, ‘under’ is an agent, therefore some

⁷*Ibid.*, pg. 9. The example is quoted by Borer from HARLEY H. & NOYER R., ‘Mixed nominalizations, object shift and short verb movement in English,’ *Proceedings of the North Eastern Linguistic Society* 28 (1998), University of Massachusetts at Amherst: GLSA.

kind of creature. Besides this kind of (syntactic) determination of the ontological role of what the term may stand for, our encyclopedic knowledge about the world provides a possible account of what might have that role. In the concrete case, while the syntactic context of (2) determined that ‘five’ in that context has to refer to some kind of action, our encyclopedic knowledge of the word ‘five’ tells us what conceptual area ‘five’ is about. We then put together the information, picking from the vast conceptual area triggered by ‘five’ the meaning (or the meanings, as Borer suggests) corresponding to some kind of action.⁸

The distinction between the linguistic component and the conceptual one is also shown by the fact that not all overridings of meaning are permitted, but only of meaning that is not carried by linguistic means. Thus, consider the ungrammaticality of (3):⁹

- (3) a. *a lot of wine is/are many
b. *there are too much carpet in this room
c. *too much carpets

Here, the count–mass features of the nouns cannot be coerced, because they are assigned by linguistic means (use of singular or plural form in absence of determiners). It is of no importance, for what concerns (3b-c), that ‘carpet’ is usually considered a count noun. This information, simply, is not read by the grammatical system, which decides grammaticality or ungrammaticality of sentences and phrases. This is what Borer thinks of the question, then:

We opt to leave within the linguistic computational system precisely those features which the grammar cannot override or ignore, relegating those which the grammar does regularly override to a different cognitive component, the conceptual one. If this is indeed the division of labour, the coercibility of the (ontological) mass–count distinction with respect to N-stems emerges quite simply from the fact that the grammar performs computations based on those formal features which it recognizes, and the fact that N-stems, as such, have none. The grammar does associate distinct structures with mass and count interpretations, but it does so regardless of the conceptual properties of the N-stems embedded within these structures. Coercion, then, is but the conflict that emerges when the grammar returns a computation which is not fully compatible with the conceptual properties of listemes embedded within these structures.¹⁰

⁸Cf. *ibid.*, pgs. 11 f.

⁹*Ibid.*, pg. 9.

¹⁰*Ibid.*, pg. 106.

Thus, the idea is that information about meaning is stored in two different ways. There is linguistic information, carried by computational structures, and encyclopedic information, that is carried by some concept associated with a word or a phrase. It is not the case, though, that the two systems might carry the same information. In fact, the very system in which the information is stored changes the information, for instance, with regard to the possibility of overriding it by coercion.

5.2.2 On the Relationship between the Linguistic and Conceptual Systems

As conceptual information is systematically overridden by grammar, we cannot assume any influence from the conceptual system to the linguistic one. In any case of conflict, it is always the grammar that prevails – to such a point that, e.g., in (2), ‘under’ is by no means a preposition, but a noun:

In the event of a mismatch [between the conceptual and the grammatical assignments of features], the grammar will always prevail. The interpretation put forth by the conceptual component can and will stretch, as much as possible within the confines of the concept under consideration, so as to match the rigid, absolute interpretational constraints circumscribed by the grammar.¹¹

As for the opposite relation, we might expect that, since grammar is the prevailing system of the two in case of a mismatch of information, it also has some influence on the conceptual system. Actually, it is not so. The very fact that a mismatch is possible, in fact, shows that the two systems are largely independent, and they come to a comparison only in the case of establishing the meaning of a sentence or phrase.¹² But, of course, this is not all that the grammatical system does: there is much grammar that is not semantic rules. And, for sure, we exploit our conceptual system in many tasks that are not linguistic.

The idea of an influence of language on the conceptual system¹³ may come from the fact that the concepts associated with words in different languages rarely appear to overlap. For instance, the Welsh terms *gwirrd*, *glas* and *llurjd* cover together the range of the English *green*, *blue*, *gray* and *brown*, so that while speakers of one language would speak of different hues of the same color, speakers of the other would speak of different colors. Hence, we might think that the Welsh conceptual system divides in three colors the range that the

¹¹ *Ibid.*, pg. 11.

¹² Actually, Borer assumes that such comparison takes place in yet another cognitive component, which she dubs the “‘making sense’ component” (*ibid.*).

¹³ For different positions on such an influence, and on the more general relationship between language and concepts, see § 5.3 below.

English conceptual system divides in four, and the only way to explain this difference is to appeal to the difference in the two languages. Another instance: Quine mentioned a people in whose language there is one single word for both half-brothers and pelicans.¹⁴

To assume speakers of different languages to live in different conceptual worlds may lead to paradoxes. According to Donald Davidson, we cannot reasonably claim that someone has a different conceptual scheme than ours.¹⁵ Such a claim, in fact, would entail either that our conceptual scheme and the other are absolutely different, or that there are partial overlappings. In the first case, we must note that it becomes impossible to interpret the other person's language, since, by assumption, we are not able to project our cognitive system (e.g., our beliefs) by means of the principle of charity. The principle of charity, however, is the only way we can attribute humanity, rationality, and a language to other creatures. For, what would the point be in assuming the existence of languages whose relation to the world is not recoverable by any means? Assuming an absolutely different conceptual scheme, we cannot be sure of what attitude other than holding true is crucial to communication. And even assuming that truth is still the central notion can be of no help:

And the criterion of a conceptual scheme different from our own now becomes: largely true but not translatable. The question whether this is a useful criterion is just the question how well we understand the notion of truth, as applied to language, independent of the notion of translation. The answer is, I think, that we do not understand it independently at all.¹⁶

If interpretation cannot even begin, there is no chance to be able to understand the grammar of the allegedly different language. But, without grammar, we are not even allowed to speak of language at all. It might be objected that there can be a grammar, even if we are not able to understand its rules. Again, let me stress that there is no point, from a scientific point of view, in assumptions of existence of entities or laws that, by stipulation, we cannot test.

Now we can see also why even assuming some partial overlapping between conceptual schemes cannot work as an example of different conceptual schemes. Such overlappings, in fact, must be necessarily construable and overrideable, and this can be done only by assuming that differences are small, contingent, and periferic with respect to the whole system of beliefs:

Since charity is not an option, but a condition of having a workable theory, it is meaningless to suggest that we might fall into massive

¹⁴Quine (1960), pg. 77; his source is LIENHARDT G., 'Modes of Thought,' in EVANS-PRITCHARD E.E. (ed.), *The Institutions of Primitive Society*, Oxford 1954.

¹⁵Davidson (2001), chpt. 13.

¹⁶*Ibid.*, pg. 194. Recall, in reading this passage, that in Davidson's view translation and interpretation are fundamentally the same task.

error by endorsing it. Until we have successfully established a systematic correlation of sentences held true with sentences held true, there are no mistakes to make. Charity is forced on us; whether we like it or not, if we want to understand others, we must count them right in most matters. If we can produce a theory that reconciles charity and the formal conditions for a theory, we have done all that could be done to ensure communication. Nothing more is possible, and nothing more is needed.¹⁷

Thus, we have seen that there is no theoretic way to associate the differences we witness among languages to differences in the speakers' conceptual schemes. Fortunately, we need not make such an association. We can simply think that the mapping of words on concepts is different in different languages, and still believe that the concepts are the same for every human being.

5.2.3 The Linguistic System

Now that we have sketched the cognitive topography of the theory under examination, let us consider in somewhat greater detail the two systems involved in it. Let us start with the linguistic system. We must ask some questions about it: first of all, whether it is a holistic system, and how we can tell.

Then, even if we assume independence between this system and the conceptual, we can ask whether the linguistic system is eminently human or a theory describing it might be considered more general than a theory about merely human linguistical processing. In other words, we must show in what sense we consider such a theory of the linguistic system a formal theory.

Finally, we must account for the differences that we see between linguistic systems, keeping in mind that radical differences give rise, under the present perspective, to impossibility of interpretation. Of course, this is not a welcome result: it clashes both with evidence (we know that we can interpret every sentence in every language in the world, no matter how different from our own) and with our initial assumptions (in fact, we have stressed several times that only if the task of interpretation is possible we are allowed of talking of language at all).

Structural Holism

To approach the answer to the question about holism, let us consider what the picture of the linguistic system as it emerges from the theory sketched in the past sections is. What we have is that, mirroring the separation between linguistic and conceptual system, lexical material in a language can be divided into two different categories: those elements which are hints of

¹⁷ *Ibid.*, pg. 197.

the compositional structure of the sentence – call them **functional elements** (e.g., articles, quantifiers, numerals, and also inflectional morphology)¹⁸ –, and properly lexical items, which Borer calls **listemes**¹⁹ and considers capable of being coerced by compositional structure in the sense given to this notion in the discussion above.²⁰

If we consider again sentences like (2), we can see that the division sketched in the last paragraph is not as sharp as we may think at first sight. Numerals and prepositions are certainly involved in the compositional structure of language: therefore, they should count as functional elements in the proposed distinctions. However, they can be coerced in other roles as well.

Two explanation of the phenomenon can be given. One is that, after all, what looked like a feasible dividing line corresponds to no actual distinction. But such an explanation fails, in fact, to explain the phenomenon of coercion at all. To see that this is true, consider once again (2):

- (2) The red under fived lunch.

It is definitely true that here *under* and *five* (which would normally count as functional elements) are coerced, respectively, as a noun and as a verb. But it is also undeniable that this coercion is operated because we can reconstruct a syntactic form in which they have such roles by means of other functional elements, viz., *the* and the past tense attached to *five*.

We might think that this shows that, after all, the distinction between functional elements and listemes is a real one, but numerals and prepositions belong to the latter category, not to the former. However, consider (4) and (5):

- (4) a. Five wines.
b. *Five wine.
- (5) a. Under pressure.
b. Under good.

It cannot be denied that here both *five* and *under* have the functional role required to assign respectively *wines* and *pressure* to the role they have in the expressions. This is proved also from the fact that it is because of the functional properties of *five* that (4b) is ungrammatical;²¹ and it is because

¹⁸Borer distinguishes two kinds of functional elements: head features, which combine with lexical items though not morphological in nature (the plural inflection of English being an example of these), and “independent grammatical functional formatives,” which can be bound or free, thus combining or not with lexical items (the morpheme *-ing* being an example of the former and the article *the* of the latter). Cf. Borer (2005a), pgs. 31 ff.

¹⁹*Ibid.*, pg. 11.

²⁰*Ibid.*, pg. 10.

²¹Of course, now that we have learned how to treat *five* as a verb, we might assume that (4b) contains it as a verb, for instance in an imperative form; thus, it would be deemed grammatical. I would not consider this case, now, because what I am interested in is why

of the functional role of *under* that we can say that *good* in (5b) is a name rather than an adjective.

It seems, thus, that some items can belong to both the categories we have introduced, or, which is probably the same, that the division we want to trace is not sharp. Nonetheless, there is a distinction to be made, as it is proved from the fact that even in the case of words that can behave both as functional items and as listemes we are able to distinguish very well the former cases from the latter ones.

Therefore, the explanation of the dividing line I am proposing goes along the lines drawn by Borer:

[...] it appears desirable to assume that it is the structure which determines the category membership of such items [as *five* or *under*], classifying them as functional or lexical, in the relevant sense.²²

The fact that such an intermediate “twilight zone”²³ between the two categories exists, and that the categorization of the elements in this twilight zone is decided by means of syntactic context shows that even a simple information like this – i.e., whether an element is functional or lexical – is not atomistically stored. This does not hold for all items: some are completely functional in every use, like articles or demonstratives (except for use in quotational environments, of course).

It is important to stress, however, that these fully functional elements cannot be the base for an atomistic semantics. In fact, their contribution to the global meaning of the sentences they appear in is, if we consider them as lexical items, very small: they become significant only when we associate to them a structure (both syntactic and morphological) that they can impose on other lexical items to coerce an interpretation of their role in the sentence.

We may think that this structure is part of the lexical meaning of these functional items. However, it is not at all clear that this move roots meaning atomistically. In fact, what the fully functional elements actually do is instan-

the phrase can be judged ungrammatical when it is, not whether there are chances that is not.

²²*Ibid.*, pg. 10 n. 4. Consider, e.g., the following sentence, as opposed to the double reading of (4b) I hinted to in the previous note:

- (i) a. *I tried five wine
b. I tried five wines.
- (ii) a. I tried fine wine.
b. I tried fine wines.

In (i) it is undeniable that the grammaticality or ungrammaticality is given by the whole structure of the sentences, and that *five* has in both sentences a functional role that helps us to discriminate between them; such a functional role, on its turn, is bestowed to *five* from the structures of the sentences. Contrast this with the non-functional role of *fine* in (ii).

²³*Ibid.*

tiate the rule of the language, but only some syntactic rule each. The meaning of a sentence is, of course, strictly dependent on the rules employed in it, but rules can build up a meaning only if considered together. The contribution of one single rule to the meaning of a sentence is dependent on the other rules that are used in the sentence. And which rules are to be used in a sentence to give rise to a particular meaning depends on the whole system of rules.

Instantiation

Let us now skip to the second question mentioned above: whether the system we have been sketching in this section needs to be instantiated by the human mind or the human brain or not. Let us start from Borer's opinion on the matter:

Within any generative approach to the study of the language faculty, what is common to all languages may, in principle, come for free. While it might still be in need of description, the null hypothesis is that it is fundamentally part of the human biological structure, however described. The true challenge is to account, within a biological approach to language, for precisely those facets of linguistic behaviour which differ from one language to the other, and hence, by assumption, are sensitive to the facets of the input.²⁴

We note that the question about the nature of the faculty of language is strictly related with the question about linguistic variation, which will be the subject of next subsection. We note further that the hypothesis of a biological instantiation of the language faculty is put forward to account for the universal characteristics shown by all human languages. However, nothing here contradicts the picture that Goldsmith has in mind, which was the subject of § 5.1 above. In effect, we can think that language has some universal characteristics and that these characteristics are biologically instantiated in human beings, or even that they are consequences of human biology and, nonetheless, think that the right way to investigate these characteristic is not through appeal to the biological mechanisms that give rise to them. This is, indeed, what modern linguistics does. Moreover, the fact that some peculiar characteristic has a determined (evolutionary) origin does not preclude the possibility that different (evolutionary) processes may result in the same characteristic. Speaking of biology, this happens every time: the eye, for instance, is said to have been "invented" several independent times.

The very separation between the linguistic and the conceptual spheres shows that there is no reference, here, to human beings. In fact, the theory amounts to saying that the compositional part of our language is independent

²⁴ *Ibid.*, pg. 261.

from our concepts. This is the same as claiming that there may exist creatures with different concepts but the same grammar as our own.²⁵ Or there could be some entity, for instance some machinery, which instantiates both our conceptual and linguistic systems, but in a very different way: viz., not biologically.

It is only when we want to confine our linguistic investigation to human language – i.e., to the language of human beings as it is different from, e.g., the language of any machine or any other creature which may ever be capable of linguistic capacities – that we need to furnish an explanation of the instantiation of the linguistic system in human beings. Such an explanation, though, cannot be a task for linguistics. Rather, it is an issue that psychology or neurology should address.

Linguistic Variation

When we ask for an explanation of linguistic variation, we can receive two different kinds of answers. In general, asking for the origin of a phenomenon can be interpreted as asking for the genetic origin – i.e., what fact in the past is responsible for the phenomenon – or for the mechanism that originates the phenomenon. Of course the two readings are connected: after all the mechanisms must have a genetic origin themselves, and we can count that as the genetic origin of the phenomenon altogether.

To answer to the first kind of question calls for a theory of acquisition, which has to explain how all human beings acquire different languages in the first years of their lives. This is a still debated subject, and a lot of literature has been written on it. We need not summarize it here.

As for the second reading of the question, we are asking for a theory to explain what makes languages different. It is obvious that different languages have different rules, and that is what makes them different. But our question now is what the difference between rules amounts to.

The standard and most widely accepted answer to this question has been based for decades on the Chomskyan Principles and Parameters approach. In a nutshell, this theory claims that every human being is born endowed of a Universal Grammar (UG), which sets several linguistic universal to be equal for all human beings (and which is sometimes claimed to be biological, although this claim might be seen as not necessary), while leaving some underspecified rules. These, in turn, are set by the process of linguistic acquisition. Rules are underspecified, but they are not totally unspecified: the freedom of the system is between the choice of one or another settings of precise parameters, which are to be thought of as switches capable of a small number of fixed positions, possibly only two positions for each switch. Within this approach,

²⁵We have given independent arguments to think that such a thing as having different concepts is possible only within certain limits. See § 5.2.2.

research has flourished to individuate parameters – possibly a small number of them – and to establish a reliable hierarchy among them.

Note that, within this theory, together with a common core of rules shared by all languages, it is assumed that all languages are somewhat different and apart from each other depending on how many parameters received a different setting. This should explain the difficulty of learning a new language after the time of acquisition, and the fact that some languages can be easier to learn in the adult age than others.

We ought to be explicit, though, on the strength that we attribute to such settings. For, if they are too strong, it is hard to see how they can be overridden in the learning of a new language. And if they were not strong enough, we could not grasp the rules of our own language.

Possibly, everything depends on acquisition, i.e., on the way parameters are set, and on what hints we have to set them. Under the view that we have sketched here, the natural candidates for the role of hints which trigger the acquisition of rules are the grammatical formatives, i.e., the functional elements of the sentences. If this is true, then functional elements are also the ultimately responsible for linguistic variation.²⁶

Also, if this view is correct, a partial revision of the Principles and Parameters approach is needed:

[...] rather than looking at a picture of language variation which includes ‘big parameters’ which involve disjunctive statements of principles or the domain of their application [...], we are attributing the distinct properties of any single E-language to whatever emerges from the interaction between universal principles and the specific subset of computational properties associated with members of the functional lexicon. As such, the set of constructions instantiated in any given language is determined by the subset of UG-specified features which are associated with the inventory of the specific formatives available in any given language.²⁷

Again, if this is correct, it means that all languages are fundamentally one and the same language, and they differ only because of their different lexicons, and of the different rules associated to parts of lexicon. I would like to quote another passage by Borer to state what is the consequence of that:

It is worthwhile noting that if we are indeed on the right track, the expectation that emerges is that there should not be any syntactic variations between languages which are substantially different from variations which may be found language internally. To the extent

²⁶Cf. Borer (2005a), pg. 264.

²⁷*Ibid.*, pg. 261. “E-language” stands for “external language,” i.e., the language as it phenomenically appears, as opposed to the I-language (internal language), i.e., the language as a system of rules, as a grammar.

that some languages appear to have a clustering of properties, these should be reducible to morpho-phonological generalizations which may constrain these languages. [...] In turn, should it turn out to be the case, as is often true, that the morpho-phonological generalizations under consideration are not iron clad, we expect, and indeed we find, a given language instantiating more than one mode of range assignment to a particular value.²⁸

Interpretation Again

Note, finally, what relationship such a theory bears with the issue we examined in last chapter, viz., interpretation. We might think that if we can assure ourselves of the soundness of the present approach, e.g., by testing its empirical consequences, then we may have a proof for the possibility of radical interpretation and for some attenuated form of the principle of charity. For, if every language can potentially express all meanings that other languages do *by means of the same rules* as those other languages, it is easy to see that we could figure out the meaning of every sentence we hear by appealing to our cognition of our own grammar and the means it provides us with.

However, consider that the possibility of radical interpretation is a necessary premise of this scientific enterprise, and cannot therefore be proved by its results. First of all, without that premise it would be difficult to identify languages, which are the very subject of study. In fact, we have seen that only to systems susceptible of radical interpretation by means of the principle of charity is appropriate to refer as languages. Thus, the principle of charity and the radical interpretation that it enhances have more power than the present linguistic theory. The latter, as all scientific theories, puts forward an explanation of some phenomenon, but the way the phenomenon is known and interpreted before the theory can work on it cannot come from within the theory.

The phenomenon to understand, here, is the possibility of interpreting all languages. If the linguistic theory came first, then the possibility of interpretation would not be assured, and the possibility of an error would always be open. The error I am referring to, though, is not the kind of error that a trial and error procedure could individuate and take care of, nothing that a principle of charity whatsoever could account of. Without some metaphysical premise for what concerns the common nature of all languages, we could never know that the universal principles that we are finding really apply to all languages. But it is exactly this premise that we are applying when discharging some theory that makes untenable claims about the nature of the universal principles: if we were not assured of the existence of universal principles, we could not even wonder whether those that we hold for universal principles are really such.

²⁸ *Ibid.*, pg. 265.

I think that here is exactly the same thing as with every other scientific theory. All science, in fact, has some metaphysical premise – e.g., the nomological homogeneity of the universe across time, i.e., the assumption that scientific laws do not change from time to time. Such premises cannot be proved by any scientific discovery, because they are the very origin of every scientific discovery. Thus, discoveries by physicists cannot be used as proofs of the nomological homogeneity of the universe across time.

Of course, the success of some scientific discipline depending on some metaphysical assumption can be used as an indirect proof of the plausibility of that assumption. It can still be maintained that this tells us more about ourselves and our scientific practices than about the metaphysical constitution of the world. This division, though, may be untenable.²⁹

5.2.4 The Conceptual System

Let us now concentrate on the conceptual system, as it emerges from this perspective. First of all, let us recall what the function of such system is, according to the present view. While the linguistic system sets the categories of the terms appearing in a sentence together with the roles that the sentence attributes to their references in the event as it is described by the sentence itself, the conceptual system must connect this skeletal structure to the world, providing references to the terms.

The conceptual properties of listemes are twofold:

It seemed to me that to be able to use a word is, on the one hand, to have access to a network of connections between that word and other words and linguistic expressions: it is to know that cats are animals, that in order to arrive somewhere one has to move, that an illness is something one may be cured of, and so forth. On the other hand, to be able to use a word is to know how to map lexical items onto the real world, that is, to be capable of both *naming* (selecting the right word in response to a given object or circumstance) and *application* (selecting the right object or circumstance in response to a given word). The two abilities are, to a large extent, independent of each other [...] The former ability can be called *inferential*, for it underlies our inferential performances (such as, for example, interpreting a general regulation concerning animals as applying to cats); the latter may be called *referential*.³⁰

Thus, semantics is claimed to be, overall, threefold: there is a referential component, an inferential component, and a compositional or structural component,³¹ which properly pertains, as we have seen, to the linguistic system.

²⁹Cf. Davidson (2001), chpt. 13.

³⁰Marconi (1997), pg. 2, emphasis in text.

³¹Cf. *ibid.*, pg. 77.

The last kind of properties, as we have seen, are holistic. The inferential component seems to be holistic too, because inferential properties can be identified with the role that the expression play in inferences, and there is no limit to the number of inferences an expression can play a role in.

Fodor and Lepore ask how inferential roles are to be identified: for, depending on their being more or less fine-grained, different accounts of synonymy can arise.³² But there is no need to identify inferential roles once and for all: in fact, our intuitions on synonymy and coextension, if any, can vary across time, and we would want our semantic system to allow for such variation.

For what concerns referential properties, consider that it is true that the referential relation necessarily goes from one symbol to the world, and that there does not seem to be any appeal to the symbol position in the linguistic system. But it is also true that such relations are acquired all together. A referential system is like a classification of perception into categories, and in such a classification every category is defined by its boundaries, i.e., by its relations to other categories.

Fodor and Lepore object that the independence between the two kinds of properties cannot prevent “radical mismatches between intension and extension.”³³ But this is good, because we indeed find such radical mismatches in brain-damaged patients.³⁴ Our question could then be why such mismatches do not appear often. “Water” could have “*the extension of a kind term but the logic of a number term.*”³⁵ I think that an answer to this point should make reference to the way expressions are acquired and used. In other words, it is a matter of enhancing communication that such mismatches do not appear.

I will not enter into details about the reasons that led Marconi into thinking that the referential and the inferential competence are largely independent one from another. Suffice it to say that there are both empirical evidence and philosophical arguments sustaining this claim. Granted this, let us concentrate with some more details on the two subsystems we have individuated.

Inferential Competence

From the examples given above in the passage by Marconi, we can say that knowledge concerning the inferential aspect of our concepts is what we associate with the phrase “world knowledge.” The problem with this kind of competence is that it is not at all clear how to distinguish the relevant inferences from the irrelevant ones. Consider (6):

(6) John has a moustache.

Provided that I have a moustache, too, both (7) and (8) are valid inferences.

³²Fodor and Lepore (1992), pg. 169.

³³*Ibid.*, pg. 171.

³⁴Marconi (1997), chpt. 3.

³⁵Fodor and Lepore (1992), pg. 171, emphasis in text.

- (7) John has hair on his upper lip.
(8) John and I have at least one facial feature in common.

One would say that to be able to infer (7) is necessary to be considered competent with ‘moustache,’ while to be able to infer (8) is not.³⁶ We can note that the difference between (7) and (8) is that the former is what has been called an analytic inference, since it is grounded on the meaning of the word ‘moustache,’ and the latter is a synthetic inference – it is grounded on contingent world knowledge.³⁷

The problem, in a slightly different formulation, is that inferential roles seem not to be compositional: e.g., we can believe that neither “brown” nor “cow” entail “dangerous,” while thinking that “brown cow” does entail “dangerous.”³⁸ The problem is that every time we try to make adjustments like taking “cow \rightarrow x such that, if brown, then dangerous,”³⁹ we will end up stating that all inferences are analytic.

I do not think that such conclusion is necessary, nor that it is a real problem for us, because if we reject the a/s distinction, we reject also its terminology. Remember also that the notion of analyticity appeals to some notion of meaning, which is precisely what it should help to define.⁴⁰ We should look at the question on the other way around: despite Fodor and Lepore tell us that thinking that all inferences are analytic is “perfectly mad,”⁴¹ the theory we are sketching is precisely that (at least a part of) meaning is to be identified with inferential role; therefore, all inferences are constitutive of meaning, and this is exactly to say that all inferences are analytic. This may simply mean that we do not have a clear intuitive idea of what is analytic.

On the other hand, it seems obvious that some inferences are less important than others. Perhaps this is true. But this difference in importance emerges only when we compare the meanings of our words with those of other people, and find that we associate quite different inferential roles to the same expressions. In other words, there is not a constitutive problem about thinking that all inferences are analytic (which is the same of depriving analyticity of any use, of course, since analyticity is a notion that has sense only if opposed to a notion of syntheticity), but only an interpretive problem, because, were it true, we could not share another person’s language if we do not make the same inferences. But this is fine: a theory of interpretation will account for our successful linguistic interactions even in the absence of a common language.

³⁶Cf. *ibid.*, pg. 27.

³⁷Cf. § 2.1.3.

³⁸Fodor and Lepore (1992), pg. 177.

³⁹*Ibid.*, pg. 181.

⁴⁰Cf. Quine (1953), chpt. 2 § 4.

⁴¹Fodor and Lepore (1992), pg. 182.

Of course we might not want to jump to the conclusion that any distinction between analytic and synthetic truth is completely useless.⁴² It may certainly have some practical use: when we are faced with contradictions that force us to revise some of the truth values we associate to our sentences, there are some immediate candidates for such revision, while other sentences seem more difficult to consider false. We can say that the former are the synthetic sentences and the latter are the analytic ones. However, there is nothing so certain about which sentences cannot be revised. We can discover, for instance, that we have been misusing a word all the time, and that its meaning is different from what we thought it to be.

My conclusion on this issue is that there is some intuitive value in the distinction between analytic and synthetic sentences, when we intend such distinction as a hierarchy of sentences according to their availability for revision. Such hierarchy, though, cannot have any use in our inquiry about meaning, for it is not definitive at all. Hence, we need a picture of inferential competence that avoids mentioning analyticity, even surrogate hierarchic analyticity.

Individual Inferential Competence Whether such an account of inferential competence (one which avoids bringing back the analytic/synthetic distinction, and respects our intuitions regarding the different relevance of (7) and (8)) is possible depends on which picture of language we have in mind. Our theory would be very different if we think that what we are looking for is meaning in a common language – i.e., in a language that is shared among speakers – or meaning in idiolects.

We have already seen that it is hard to sustain a classification of common languages without considering idiolects primitive. In other words, the notion of common language cannot be but derivative. Let us consider what is the relationship between individual competence and a linguistic community, in Marconi's opinion:

[...] at the lexical level, individual competence, even idealized individual competence, does not coincide with encyclopedic knowledge, that is, with the totality of true belief that can be ascribed to the linguistic community as a collective entity. We also saw that is extremely hard to isolate a plausible subset of encyclopedic knowledge as being constitutive of (lexical) semantic competence, that is, a set of propositions we all ought to know or believe in order to be regarded as lexically competent. As a matter of fact, each of us knows or believes partly different things (partly different subsets of the collective encyclopedia), yet we are all competent in the use of our language. Could lexical competence be defined

⁴²See Marconi (1997), pgs. 32 ff. for some discussion about the possibility of reviving some humble concept of analyticity.

as the *intersection* of such different individual competences? In principle, no, for we have no guarantee that the intersection of all competences is rich enough to be considered adequate as a competence (it is quite possible that the set of beliefs we all share about, say, gold is very small, too small to constitute an adequate competence with ‘gold’). It remains that we regard the notion of absolute competence as empty at the lexical level: on this level, we only have *competences*, in the plural. We ought to speak not of a unique lexical competence, only of individual competences. In this sense, there is no language, only idiolects.⁴³

Our problem of distinguishing between a theory of competence for idiolects and a theory of competence for speakers in a linguistic community can be reframed in another dicotomy, between a theory of individual competence and a theory of what Marconi calls **convergence**,⁴⁴ i.e., a theory to explain how individual competences meet in the ordinary use of language as a means of social interaction. The latter is equivalent, in my opinion, to a theory of interpretation, and the exposition given in the first part of this chapter can account very well for it.

If what we are interested in is individual competence, the distinction between (7) and (8) cannot entail any problem. Let us consider again (6). We want to know if Mary is competent with the word ‘moustache’ and ask her if (7) and (8) are true inferences, given (6). We expect her to say that (7) is a correct inference from (6). As for (8), we expect her to say that, if she knows me, it is a correct inference from (6), or, if she does not know me, that the inference would be correct if I had a moustache myself. In other words, the alleged analytic or synthetic character of a given inference is not relevant to decide of someone’s competence. What is relevant is Mary’s being prepared to draw both (7) and (8) as inferences from (6), given some background knowledge or experience. And there cannot be any limit to the inferences that she should be prepared to draw, if she is competent with ‘moustache.’

A moral that we can draw from this story is that appeals to our intuitions are always disputable, because our intuitions themselves are not as clear as we would like them to be. Besides intuitions about the dependence of (7) on word meaning, in fact, there are intuitions about the fact that all that we say depends on word meaning.⁴⁵ Must we decide between some analytic/synthetic distinction (or, better, word meaning/world knowledge distinction)⁴⁶ and the intuition that every inference depends on word meaning?

⁴³ *Ibid.*, pgs. 52 f., emphasis in text.

⁴⁴ *Ibid.*, pg. 172 n. 39. Cf. also *ibid.*, chpt. 4.

⁴⁵ Such intuitions can be so strong that we want to call it a truism.

⁴⁶ If intuitions about a distinction between what in our sentences is due to word meaning and what to world knowledge is perfectly maintainable, it seems very unpalatable that there exist intuitions about the analytic/synthetic distinction.

I think we need not choose. As long as we need a guide for everyday use of language, our intuitions about it are the best guide we can imagine, albeit they might be blurred or even inconsistent one with respect to the other. After all, our intuitions have been working well ever since human beings acquired their language faculty. As for philosophical or scientific reasoning on the subject of language, all intuitions must yield to evidence and arguments.

Coming to the point in question, we are not claiming that all inferences are a matter of word meaning. We are simply stating that there is no way to tell when it is a matter of word meaning and when world knowledge enters the picture. Maybe this distinction is simply too blurred to be of any use for the present discussion. Moreover, and partially because of this fact, there is no way to tell how much of the conceptual system is needed to be assured about the competence on one single concept. One thing is for sure, though, and I will repeat it: there cannot be any limit to the inferences that one must be prepared to draw, to be considered competent with a concept.

Inferential Holism The question now is whether this absence of limit amounts to holism. Marconi thinks that the answer is no, because holism has unwanted consequences for a theory of competence:

Now suppose that holism holds, so that a constituent's semantic value (like the semantic value of every expression of any language) is a function of the linguistic system as a whole, in any *prima facie* plausible sense: for example, a function of the system's structural properties *and* the semantic values of the system's other units. It follows that understanding any sentence involves the understander's entire semantic competence, for it requires that information be used relative to the whole linguistic system, structure and elements. Understanding would thus be impossible for a mind like the human mind in memory and computational resources. Only God would understand the sentence 'All men are mortal'.⁴⁷

Marconi's position, then, amounts to a form of molecularism, because he claims that some subset of the whole system can do, although he admits that there is no way to establish which subset:

As I anticipated, the moral is the following: we feel that our knowledge of how to use a word is inextricably intertwined with our knowledge of how to use many other words. This is true up to a point: it does not entail that to be able to use one word, we must know how to use all other words, nor that we must know how to

⁴⁷ *Ibid.*, pg. 49, emphasis in text.

use the other words to the same extent as we know how to use the original word.⁴⁸

I would like to note that Marconi, here, is assuming again a communitarian perspective on language. His mention of all the words, and also of the extent of knowledge about the use of a word, makes me think that what he has in mind is evaluating such an extent with respect to a norm of use established in a community. For an analysis of the problem from the point of view of idiolects, though, it is not a relevant objection that one needs not know how to use all the words used by the linguistic community. Our problem is, in fact, whether the global knowledge of how to use all the words that one knows how to use is involved in the knowledge of how to use one single word. This view does not entail that only God would understand ‘All men are mortal.’ Such a conclusion is entailed, in fact, by the former view, which Marconi rightly refuses, and which, as we have seen, is different from what semantic holism about idiolects would claim.

Again, let me stress what the problems with a molecularistic view are. If you think that some subsystem is sufficient to assure inferential competence with a certain expression, you have two options as for the identification of such subsystem. Either you believe that we cannot identify such subsystem and we will never be able to, or you believe that such an impossibility is just due to defects of the theory, which will be mended with further research.

The former position strikes me as dogmatic and undefensible. The latter is at least coherent, although there is a lot of job to do in order to shape it into a working theory. We need an effective strategy to individuate the relevant molecule attached to a certain expression, and we need it not to be a holistic strategy. Until now, there is no sign of such strategy in sight.⁴⁹

Referential Competence

Let us now take a look to the referential component of the conceptual system. It seems that, at least for what concerns this, holism is out of the question. Indeed, it seems that an explanation of the connection of each concept with its reference (i.e., with some part of the world or the other) can be explained without appealing to the connections of other concepts. A causal connection, for instance, could explain reference.

The idea behind such a suggestion is that the connection between the conceptual system and the world is operated by the perceptual causality of the world on the conceptual system. But perception is not enough to assure the wide variety of features that distinguishes one concept from the other. The problem is that the causal theory of mental content is, fundamentally, an empiricist theory, and, as all empiricism, has to confront with the rejection of

⁴⁸ *Ibid.*, pg. 55.

⁴⁹ On this subject, cf. § 3.3.1 above.

what, according to Quine, is the second dogma of empiricism, viz., the distinction between observational content and theoretical content. In particular, the rejection of the dogma is at the basis of the claim that theoretical – i.e., conceptual – content is always involved in any observation.⁵⁰

One way to read such a claim is to see every observation as in some way already interpreted. Another way to put it is to claim that there is no such a thing as mere observation, as long as what we are interested in is the relation with the conceptual system. Of course, objects (in a broad sense of the word) have a causal effect on the conceptual system; but in order to do so, they must be filtered by some cognitive structure. The purpose of such cognitive structure should be to make the referential system be like a classification of perception in categories. As we already pointed out, in such a classification every category is defined by its neighborhood, i.e., by its relations to other categories, and this is how holism is entailed by this picture.

One reason to believe that there cannot exist such a thing as pure observation is that, according to cognitive studies, not even perception can be deemed the pure reception of external stimuli. Every perception, i.e., every perceptive stimulus that plays a cognitive role, has precisely the cognitive role it has in virtue of the possible interaction between the subject and the object of the perception. Such a possible interaction, when recognized in such a way by the subject, is called an **affordance**.⁵¹

Perceptions, with their affordances, are not yet to be considered concepts. We should ask what is the difference between the two notions. The major difference is that concepts are modality free, in the sense that they are not necessarily connected with one sensory modality, as perceptions inevitably are. In spite of this, it is not at all evident that concepts lose their affordances also. It seems, indeed, that an important part of the way we conceptualize entities is knowledge of the kind of relation – therefore, of interaction – that they can bear with ourselves. In fact, we cannot have the concept of something which we do not know how to relate to, or at least how to compare our position in the world with its own. Even the most abstract concepts we could entertain are clearly there to be exploited by abstract reasoning.

Ontological issues are not at stake here: we can have a concept of some entity even if we are not quite certain of what kind of entity it is, or if it is a legitimate entity at all. What matters, in terms of our conceptualization, is the relationship between ourselves and those putative entities, whatever their ontology should turn out to be. For instance, it may happen that some kind of abstract entities (e.g., properties) are reducible to concrete entities by way of some metaphysical theory. This reductivism does not prevent us from having concepts of properties, nor it implies that our concepts of properties reduce to our concepts of concrete objects. This is because our relationship with

⁵⁰Hanson (1958), pp. 8 ff.

⁵¹Bermúdez (1998), pgs. 112 ff.

properties is different from our relationship with concrete objects: even if they should turn out to be the same ontological kind of entities, we are nonetheless conceptualizing them in different ways. This difference in conceptualizing depends on the difference of relationship, because, accordingly, that is the only difference between those two kind of entities.

For a less controversial instance (though not at all an uncontroversial one), consider the case of tracks in a bubble chamber: where I see merely tracks, trained physicists see different kind of particles. This is because trained physicists know how to interact with particles, whereas I do not know, therefore they know how to apply the concept to the tracks in the bubble chamber, and I do not know. And this application is what constitutes the referential competence.

This example does not imply, however, that I and the physicists have different conceptual schemes: only that we have different practices, because we have different goals, and beliefs, and so on. Now, if the individuation of the reference of the concepts depends on the relationship between the subject and the world, then what is important, in this process, is to consider the subject's cognitive system as a whole. For, to understand one's relationship with the world, it is essential to regard oneself in a certain way – which is sketched in § 4.3 above. To make a long story short: we interpret our behaviour and our interaction with our world according to an axiom of rationality. By doing this, we are able to connect the concepts to their references, because only certain such connections can make our actions rational. This can entail a certain amount of underdetermination, especially intersubjectively: but this is not a problem if we consider that intersubjective connections between concepts (or language) and the world should be controlled by the principle of charity and, in general, by the mechanisms that enhance interpretation.

5.3 Cognitive Semantic Theories

Let us now examine another kind of semantic theory, according to which the linguistic side of semantic phenomena cannot be discerned from the conceptual side. We can also put it in a somewhat weaker formulation, and say that there is some interdependence between the two aspects of semantic phenomena that allows us to think that we are facing a continuum, probably originated by interactive systems. The result does not change, as we must accept the practical impossibility of making clear distinctions in the fuzzy boundary between the phenomena imputable to one or the other system.

Of course, the relation between the language faculty on one side and other cognitive faculties on the other can go both directions. I believe that this is the reason why there is yet no agreement about the extent of the influence of language on conceptual structure and on reasoning and, conversely, on the influence of conceptual structure and reasoning on language.

One of the sides of the debate is represented by what is known as the Sapir-Whorf thesis (or hypothesis). It is the hypothesis according to which the way a particular language is structured and the features of reality that a particular language is able to convey shape the culture of human beings speaking that language, by affecting the way they look at the world. Whorf made some controversial examples about Hopi language not having means to express the notion of time and, consequently, about Hopi people not having any notion of time at all.⁵² He also thought that gesturing while speaking is typical of Western culture, and is connected with the fact that, in Indo-European languages, most talk about the mind sphere exploits metaphors from the physical and spacial spheres.⁵³

Under this view, most of the semantic work is done by the linguistic system. In fact, not only the syntactic role of words in a sentence are to be reconstructed by way of linguistic rules, but also the range of possible meanings that certain words can assume is determined by linguistic rules, which give to the sentence more or less “semanticality,” i.e., make it more or less acceptable under the perspective of meaning.⁵⁴

[...] the boundary between what we formally take to be linguistic or lexical knowledge and that which is sometimes referred to as “commonsense knowledge” might appear fuzzier than ever. Once we start enriching our lexicon with information that, to a linguist, appears better suited for a knowledge base, there may appear to be no systematic means to judge where to stop [...]. Yet, the fact that there appears to be a continuum between these two types of knowledge doesn't mean that there are not clear cases of paradigmatic linguistic behavior that are better treated as language specific knowledge, rather than in terms of general inferencing mechanisms. One of the goals of the present work has been not only to argue in favor of richer lexical representations and compositional mechanisms, but also to show that these structures and devices are still language specific in identifiable ways. What is different about the approach taken here is that there are clear and obvious means to interface lexical knowledge with commonsense and pragmatic inferences.⁵⁵

In order to put conceptual information in the position of playing a linguistic role, Pustejovsky gives great importance to lexical information. This move is, as we can recall, exactly opposite to Borer's one, which amounts to the claim that all lexical information is, *qua* lexical, irrelevant to the linguistic system.

⁵²Cf. Whorf (1971), pgs. 57 ff.

⁵³*Ibid.*, pg. 154 f.

⁵⁴For the notion of “semanticality,” cf. Pustejovsky (1995), pgs. 40 ff.

⁵⁵*Ibid.*, pgs. 232 f.

In her theory, the syntactic scheme played a preeminent role, and the syntactic scheme was reconstructed by way of functional items, which, as such, were devoid of any lexical information. In Pustejovsky's theory, every lexical item carries with itself the possible syntactic schemes in which it can appear. To each lexical item, in fact, is associated a semantic structure which defines the relations between that linguistic item and all others (this is possible thanks to the individuation of conceptual lattices which define the various relations between concepts, especially hyponymy). If this skeletal grasp of Pustejovsky's theory is correct, then, in spite of what the emphasis on the lexicon and on the study of lexicality may make us expect, this can legitimately be called a holistic theory.

Let us now concentrate on the other direction of the possible interaction between the linguistic and conceptual systems, i.e., the interaction that goes from the latter to the former. The debate about this side of the issue originated as well from Sapir and Whorf's ideas. In fact, it is not clear the direction they had in mind, when they spoke about an interaction between language and concepts. The issue is of course complicated by the indistinction between language and theory, or language and culture, where the latter element of each pair stays for some kind of conceptualization of the world:

Before beginning in earnest, I should say something about my distinction between "culture" and "language." To linguists this is a natural distinction. To anthropologists it is not. My own view of the relationship is that the anthropological perspective is the most useful, but this is exactly what this paper purports to show.⁵⁶

It must be noted, in this regard, that speaking of conceptualization of the world is misleading, as we have stressed several times before. We can think of such conceptualization, though, as the way one comes to acquire an interaction with the world, and this, in a sense, is the same of acquiring a world itself.⁵⁷

According to this second view (which is more or less equivalent to functionalist approaches to linguistics), the cultural and conceptual efforts made by each human being to interact with the world have a great impact on the structuring of the linguistic system. Daniel Everett holds an extreme position in this regard. He discusses a few peculiarities of the linguistic system of Pirahã (an Amazonian language), such as the following:

[. . .] the absence of numbers of any kind or a concept of counting and of any terms for quantification, the absence of color terms, the absence of embedding, the simplest pronoun inventory known, the

⁵⁶Everett (2005), pg. 622.

⁵⁷Cf. last section's discussion of affordances. More on this issue will be discussed in the next chapter.

absence of “relative tenses,” the simplest kinship system yet documented, the absence of creation myths and fiction, the absence of any individual or collective memory of more than two generations past, the absence of drawing or other art and one of the simplest material cultures documented, and the fact that the Pirahã are monolingual after more than 200 years of regular contact with Brazilians and the Tupi-Guarani-speaking Kawahiv.⁵⁸

All these features are reconducted to one single source: “*the restriction of communication to the immediate experience of the interlocutors.*”⁵⁹ Notice the difference between this kind of approach and the one sketched before. There we had a theory which regarded the linguistic system as responsible for all possible reasonings about what words mean or, at least, for the vast set of most common reasonings about it. Here, instead, reasonings are responsible for language, but since our investigation of reasoning lies principally on linguistic clues, this means that we can safely attribute only the reasonings that a speaker can share with us.

This can have two consequences, depending on our position about interpretation. If we judge the principle of charity unavoidable, then we have to hold that all human beings reason in the same way; and since reasoning grounds language, all languages must be more or less the same, they have to share their structure. On the other hand, we may feel that the superficial diversity of natural languages is a hint of different kinds of reasoning taking place in human minds, depending on culture, that is, on historical contingencies and, perhaps, on the environment.

For a strict functionalist position, though, the first option is untenable. In fact, according to this position, the conceptual structure is the only source of linguistic change: if the conceptual structure is the same for all human beings, we have no way to explain why languages are different. Radical functionalism, hence, entails intranslatability and incommensurability between distant languages which are expression of distant cultures.⁶⁰

Radical functionalism has also other flaws.⁶¹ For instance, it is difficult if not impossible to explain linguistic acquisition from such a standpoint. The usual functionalist picture of linguistic acquisition grounds it on the typical processes of human reasoning, which, by way of some kind of induction, reconstructs the rules of the grammar. Anyway, such a general theory of human mind cannot account for the specificity of language, if some innate knowledge of what the object of such induction is is not presupposed too. In other words, the child must have means to distinguish that portion of the world which is linguistic – i.e., relevant for an induction of linguistic rules. But this is

⁵⁸ *Ibid.*, pg. 621.

⁵⁹ *Ibid.*, pg. 622, emphasis in text.

⁶⁰ Cf. Everett (2009), pg. 28 and Gordon (2004), pg. 498.

⁶¹ For a more complete examination of functionalism, cf. Newmeyer (2005), chpt. 5.

equivalent to assign the linguistic acquisition to a dedicate faculty, somewhat separate from the other cognitive faculties of the child, which is precisely the solution that functionalists do not want.

Another problem is methodological. Since functionalism must explain all linguistic phenomena relatively to the function that those phenomena have in the conceptual life of human beings, it is hard to find an explanation for rare linguistic phenomena, and in particular for those which virtually never occur spontaneously, but that can be easily elicited. In fact, strict functionalists deny scientific validity to the practice of elicitation altogether:

Each construction in each language should be studied in its appropriate linguistic and cultural contexts before drawing any conclusions about it. This is one reason I almost agree with a statement that I overheard from a field linguist a few years ago that no one should be allowed to write an article about a language before they have written an entire grammar of the language. Only after writing a grammar of a language can you contextualize your understanding of how structures and constructions and meanings work in that language.⁶²

The reason of these flaws is that a strict functionalism, i.e., a theory which, given a culture, would reconstruct a language, must deny an innate faculty of language. At most, it can accept a biological community between human beings, and explain similarity of linguistic behaviour from this community:

Whatever the *general* properties are that distinguish human brains from those of other species, those might well be investigated before coming to the conclusion that there is a module in the brain or a language organ dedicated strictly to language.⁶³

In a nutshell: functionalism is methodologically very different from other more generally accepted approaches to linguistics, and it is very difficult to think that results obtained from functionalist premises can be shared by the entire community of linguists. I would say, paradoxically enough, that strict functionalism investigates a different subject than, for instance, generativism, and this accounts for the atmosphere of incommunicability that surrounds those works where the two point of views come to a clash. However, let me

⁶²Everett (2007), pg. 12. Such a position should not be confounded with the holistic position detailed in the last chapter. The holism that here is entailed, in fact, leaves room for incommensurability and intranslatability; the holism about interpretation excludes this possibility by specifying that the comprehension of the appropriate linguistic and cultural contexts is to be achieved by projecting as much of our own linguistic and cultural background into the picture as possible. Everett, instead, seems to hint to the possibility of depriving oneself of any bias and achieving freshly uttered new linguistic and cultural structures.

⁶³*Ibid.*, pg. 28, emphasis in text.

stress once again that the two theories sketched in this section are not mutually exclusive. There is no need to endorse radical functionalism, of course, nor to adopt a theory which accounts only for one direction of the relationship between the conceptual and the linguistic systems. The two systems, in fact, can be in a reciprocal relationship of mutual interdependence. It is necessary, though, that both sides acknowledge which are the differences between the premises that they accept or do not accept, and that lead them to different conclusions. Only then they can meet on some common ground.

CHAPTER 6

Evaluating Semantic Holism

We come now to reviewing semantic holism as it has been detailed in the previous chapters, in order to point out potential flaws in the arguments leading to the present theory. Furthermore, an eye will be kept, as for semantic atomism and semantic molecularism, to the connection between semantic holism as a theory of philosophy of language and of understanding and other disciplines such as ontology, epistemology, linguistics, and so on.

6.1 Internal Problems

Although throughout the previous treatment we have already described – and (hopefully) confuted – most of the major objections which may be moved against semantic holism, there are still details for which semantic holism might be reproached. Most of the following criticisms are indeed sound, although their actual status as flaws may depend on the methodological stand from which we regard semantic problems.

6.1.1 The Problems with Interpretation

Our holistic theory of interpretation is twofold, dealing with the interpretation of other people's sentences as well with self-interpretation. Hence, in order to see whether everything is alright with such theory, we have to examine both sides of it. Although the theory of self-interpretation is arguably the most speculative and doubtable of the two, we will start with the proper theory of interpretation right away, and we will see that it is not completely unapproachable itself.

On the Management of Radical Interpretation

We have already seen what the problems with a behavioristic approach along Quine's lines are, and how we could fix them by means of Davidson's prin-

ciple of charity. We have also examined some of the objections that might be raised against such an improved theory of radical interpretation, and we managed to confute all of them. The objection I would like to consider now is quite different in its spirit, in that it does not question neither the form of the argument, nor the sustainability of the task of the radical interpretation. Rather, it points to the abstractness of the theory, to question its effectiveness for what concerns the actual practical cases of interpretation. In other words, although holistic theorists of interpretation boast that they take into consideration actual practices of interpretation, and their theories are therefore more reliable, we may doubt that it is really so.

This is an objection we have somewhat already been confronted with: semantic molecularists, in fact, claim that taking into consideration the particular circumstances of a given utterance is the only way to correctly interpret it and keep faithful to the constraints that rule the actual practice of interpretation. We have said, on this regard, that this may very well be true, but that it does not lead us anywhere if we lack an effective way of individuating relevant circumstances – or localities, as we have dubbed them after Bilgrami’s use of the term.

Our argument for the holistic theory of interpretation that stemmed from the task of radical interpretation was that, by accepting it, we put ourselves in the condition of being able to effectively restrict our interpretive practice in the way which seems more comfortable, while still being assured that, on a more abstract and theoretical level, we are moving within a general background of interpretive freedom. To put it in another way, it would be a mistake to reduce interpretation to the individuation of localities, not only because our ability of recognizing one could fail in unpredictable ways, but also because, when it does, we cannot know whence to drag into the picture the elements that would enable us to give an interpretation. Moving within a holistic picture, instead, assures us that any element is available for our interpretations, or attempts thereof.

There are three considerations still to be made on this regard. The first is that maybe the theory so far detailed is not really a holistic one: perhaps is just a more complicated form of molecularism. Instead of having several different localities which are detached and independent one from another, we can stick to a picture of several overlapping localities, sharing some element but not another, enabling some interpretations but not others.

Putting things in this way is still misleading, though. On one hand, we may still be faced with the question about the recognizability of localities. This is not an unsurmountable problem, now: it is no more difficult than implementing our theory in a trial and error procedure, and we obtain the quite welcome results that interpretations requiring less local elements are processed in less time (which we would expect), and that interpretations requiring localities already exploited will be processed in less time than interpretations

requiring novel ones.

On the other hand, the real difficulty concerns the problem of having some kind of cognition of all the possible localities that might come useful. This is arguably impossible: since we cannot put an upper bound to the different elements that are or might be necessary to interpret utterances in circumstances, we might end with having to deal with infinite many contexts of interpretation, some overlapping, some simple, some very complicated. We would have a better time if we imagine that localities – i.e., contexts of interpretation – are built, when they are needed, again by means of a trial and error procedure, which does not choose among them, but connects the various conceptual elements into one locality. Under such a picture, holism and molecularism need not be contrasting pictures. While holism gives, as I have said, the freedom we need to build all the different contexts of interpretation we might ever need, it is by means of a molecularistic theory of interpretation that we put any single one of such localities at work.

My second consideration is connected to the first – as the third will be, too. It goes like this: although everything said in the last paragraphs is very plausible, it is still a theoretical and abstract construction, and therefore it does not meet the initial objection – i.e., that holism does not give a concrete explanation of the actual practice of interpretation. In fact, it may be argued that in most cases there is no such thing as a construction of localities, since we know very well which one of many stored localities (which has worked for precious interpretations) we have to apply. Not only that, but we may even argue that we never are faced with circumstances so novel that we need to build a new context of interpretation for it.

The latter proposal seems too strong: it is certainly counterintuitive that we never encounter novel circumstances of utterance of previously known sentences, as well as novel sentences altogether. But I have expressed more than once a suspicion about appeals to intuition, and I do not think that this case is any different, at least in the relevant reading of ‘novel,’ the reading that does not make the appeal to intuition a truistic platitude. For all we know, human beings may be (genetically, biologically, or somehow ontologically) programmed for being able to interpret an infinite number of sentences and utterances in only a finite number of different circumstances. I am not going to argue against this conclusion. Rather, in what follows I will try to dismantle the premises that lead to it.

Before that, though, let me add a further consideration about the relationship between holism and molecularism that emerged from the previous paragraphs. For this purpose, let us consider once again Bilgrami’s example of the two different localities which entitle the attribution of two different water-concepts.¹ I will sum up it briefly again. My concept of water (and therefore any utterance containing the word “water” in its literal sense and

¹§ 3.2.

depending on the context for its interpretation, such as “That is water”) has a different content in the following two situations. In the first situation, I am thirsty, and, wanting to quench my thirst with the least expensive drink available, I get a glass of water from the tap. In the second situation, I am a chemist in a lab, and I am preparing a solution of one part of sulphuric acid and nine parts of water. Bilgrami’s claim is that the two water-concepts are arguably different, since the second, but not the first, has as a part of its content the proposition that water is H_2O .

I am not going to repeat here the reasons why I find such a conclusion disputable.² My point here is not a theoretical one. Even granting some theoretical possibility to Bilgrami’s example, we still have to show that things are in fact as he depicts them. And this seems to me an empirical task, something that we can assess only by speculation on the results of some practical experiment.³ For instance, we can put subjects in the two different situations detailed above (or give them a text to provide them with the same background of notions), then see how long it takes to process an utterance containing the word “water” used in its literal sense and in an utterance whose interpretation depends on the given context, once we present the subject, before the processing, with priming information⁴ which, according to Bilgrami’s account, is relevant only in one locality but not in the other, e.g., the chemical composition of water.⁵

Bilgrami’s prevision should be, in a rather uncharitable – but, hopefully, not too misleading – interpretation of his theory, that we have facilitation of processing in one case, but not in the other. Of course, this is not the only prevision that we can make, given Bilgrami’s theory, nor my own perspective on the subject entails only one possible experimental account. Rather, the question is an empirical one exactly because we need to rely also – though not only – on the interpretation of some experimental result in order to clarify our ideas about the issue.

I would like to add here a consideration of a different order about the objections I am presently dealing with. Although there is a sense in which they are perfectly legitimate, there is also a sense in which they completely miss an important point. In fact, they are dealing with semantic holism as if it were a theory about an infinite set made up of an infinite number of element, which our cognitive system must assemble in different ways in order to permit

²See § 3.3.1.

³I want to stress that I do not believe that more philosophical insights come from empirical findings than from any other source. I only want to say that practical experience, and in particular scientific experiments, are not a source that philosophy can afford to ignore.

⁴Priming information is some kind of information which the subject registers though being unaware of it, and that interferes, positively or negatively, with the cognitive task that is being studied.

⁵The first idea of such an experiment occurred to Marta Ghio during a conversation with me.

interpretation. However, this is hardly what people should understand holism to be.

Holism, as it should be clear by now, is a theory whose primitive notion is not that of a pool of elements from which to choose some and create the localities we may happen to need, overlooking the elements we do not need for the moment. The primitive notion of semantic holism, rather, is that of an organic system which has the power to signify and interpret (or to permit signification and interpretation).

The primitive notion of holism, therefore, is that of rationality. It is because our rationality is structured in the way it is that we are able to recognize into it elements and subsystems – i.e., aggregates of elements. But it is important to acknowledge that the individuation of elements and aggregates cannot be prior to the operation of rationality itself, because such an individuation depends on the operation of our rationality. As a matter of fact, the two notions tend perhaps to be coincidental: any operation of our rationality is the individuation of some kind of element or aggregate of elements that is relevant to the rational interactions of the rational being. By converse, any individuation of such elements or aggregates is, trivially, an operation of rationality.

I say that the latter is trivial, because we have already shown that any individuation of elements and aggregates is grounded on some kind of taxonomy. And taxonomy is the way rationality projects its purposes on the object of its interaction. It is in this sense, I believe, that we should understand Christopher Peacocke's idea that acknowledgment of identity is a cognitive primitive.⁶ Such an idea should not be understood to mean that there is no cognitive notion to which we can reduce acknowledgments of identities, because such a notion exists: it is rationality. It is true, though, that, even if we understand rationality to be something more than simply acknowledgment of identities, the reduction of the latter to the former is uninformative in the sense that we still lack a satisfying definition of rationality. In this sense, any reduction of some faculty to rationality as a cognitive primitive is not a proper reduction, in that it is less informative than we should expect a proper reduction to be.

Before digging some more in the rationality field, I would like to state now in what sense I think that the objections I have faced in this section are legitimate. They are legitimate because, even if they do not meet the fundamental theoretical premises of semantic holism, there is nothing that compels the opponent of holism to accept those premises. In this regard, what we have achieved by being confronted with such objections is the acknowledgement that holism and molecularism are both legitimate and coherent in themselves, although they stem from very different theoretical premises. Moreover, a domesticated form of holism can be implemented without much ado into a molecularistic account, and a domesticated form of molecularism is perhaps the best way to take shortcuts in a holistic theory of interpretation.

⁶Cf. Peacocke (2008), esp. chpt. 4.

Rationality as a Primitive

Christopher Peacocke's idea is that the recognition of identities is a cognitive primitive, and happens, therefore, thanks to an implicit conception of what is identical with a certain object (or, more broadly, with any entity which we already have fixed as reference for some of our terms or phrases). He advances this theory to account for meaning and understanding (especially the latter) in a way that bears some resemblance to the theory I have been detailing in the previous chapters. Personal experience of meaning and understanding, according to Peacocke, is crucial in order to proceed to interpretation of other people's sentences and utterances.⁷

Although Peacocke's stress is not on the globality of the experience, but rather on the single experiences which are invoked and compared from time to time, we have already seen that in order to know how to attribute some kind of mental state to other people, we have to attribute to them the general shape of our cognitive system as a whole. I think that it is quite easy to implement such a modification in Peacocke's theory, although he may not be willing to do so, for other reasons that will appear clear.

Perception and Representation of Identities The problem Peacocke has with holistic attempts to account for understanding, and the objection he would perhaps voice against an account along the lines of the previous chapters, is that the central notion, for him, is that of reference. Hence comes the stress he puts on the notion of immediate or privileged access to reference that comes from being aware of one's own mental modifications (which accounts for the semantics of mental),⁸ or from perceiving object and properties (which accounts for the semantics of external objects together with their properties).⁹

We have more than one reason to be skeptical about the latter proposal. After all, it is not at all clear that our understanding of perceptible properties comes primarily from our knowledge of what it is to be perceived as having that property. First of all, we must consider that it is quite unclear how the recognition or assumption of an identity between a perceptible property and some other property occurs. Peacocke acknowledges this, and therefore considers such a recognition or assumption, and the idea of identity that is entailed by them, an implicit conception.

While appeals to implicit conceptions can indeed be explanatory,¹⁰ this may happen only when just a few of them are invoked. In the case of identities, though, a new implicit conception has to be called for every different property which is recognized as identical with another. I will not commit myself to

⁷ *Ibid.*, pgs. 175 f.

⁸ *Ibid.*, chpts. 5 ff.

⁹ *Ibid.*, chpt. 1.

¹⁰ Cf. *ibid.*, chpt. 4, § 4.

the claim that there are infinite many such properties, but their number is definitely high (just consider, e.g., that every different hue is such a property).

Moreover, there is no sense in which recognizing that for an object to be oval it is to be as some oval object that I perceive as oval is different from recognizing that for an object to be yellow is to be as some yellow object that I perceive as yellow. Both recognitions call for one and the same cognitive explanatory principle. My claim is that this explanatory principle is rationality itself. Thus, I would stick to Peacocke's implicit conception theory, while preventing it from having too many implicit conceptions to deal with.

There is still a greater problem. Our understanding of unperceived perceptible properties must appeal to some understanding of that property that goes beyond the mere perceiving it. Otherwise any perceptual illusion would baffle not only our present understanding of the world and of the descriptions we can give of it, but also our understanding and meaning faculties altogether. Imagine, for instance, that John has never been faced with ovals, but only with open curves seen from a perspective such that they appeared oval, and that he learned to call such shapes "ovals." It is not at all clear, in my opinion, that John refers to open curves perceived as ovals when he utters the word "oval." But, under Peacocke's theory, he cannot refer properly to oval-shaped objects perceived as oval or to objects identical to oval-shaped objects perceived as oval, because he does not know at all how such objects are perceived, and he does not even know whether such objects exist at all.

Appeals to intentions and mental representations are void, because there is no detectable difference in John's talking under an intention or another, or under a representation or another, unless his hearer construes his intentions and mental representations. And these are construed by ways of interpreting John's behavior and dispositions towards, say, ovals. Among such dispositions are also his referential dispositions, the circumstances in which he is willing to attribute the property of being oval to something he may or may not perceive, and he judges more or less identical to objects he may or may not perceive. These, I repeat, are all circumstances that can be exploited as hints to construe the meaning of John's utterances of sentences containing the word "oval," but they are by no means the only sources of our understanding of them.

The problem, so far, is that an interpretation theory relying basically on the individuation of identities between references presupposes an ontology of perceptible objects and properties prior to the rational intervention of a subject's cognitive faculties. While this presupposition might seem inoffensive, the claim that our language and our cognitive system is shaped as to be in perfect agreement with such an ontology is not. Peacocke tries to justify the second claim by appealing to an evolutionary process:

My own view is that the easiest way for such complex, relationally individuated states to occur is for states of their kind to have evolved by a selection process, one which favours the occurrence of

those states whose representational content is correct.¹¹

Here Peacocke seems to misunderstand the claims of evolution theory, and what claims it can be appealed to to support. If we are undoubtedly justified in thinking that being able of representing reality as it is would be an evolutionary advantage, in the sense that it would enhance the survival of the fittest, this does not imply that our cognitive system has such a property, nor that an evolutionary story that tells how a cognitive system might acquire such a property is the easiest explanation of the origin and nature of our representational faculties. In fact, all that is required from a cognitive system as a whole to enhance the survival of the organism is that it prompts appropriate answers to given stimuli, and this may happen even if representations are constantly but coherently incorrect.

The real flaw in reasoning, though, seems to be that of thinking that there is a way in which the world is, which is prior to our understanding and representing of it, and that our presence in the world as rational perceivers and rational agents does not affect the way the world is in any way. In other words, we are probably not entitled at all in talking of the world as separate from the ideas and representations of it that we have, nor as separate from the ways we interact with it. Our struggle to make sense of the world in order to act and interact with it consigns a world to our cognitive system, so that it is meaningless to ask how the world is without a rational mind to make sense of it. I will come back again on this issue.

Another problem a conception like Peacocke's has to face is that of relations. If our understanding starts with references and truths, we must be told what it is to be counted as the reference of a relation. Moreover, we need here not the set-theoretical entity proposed by standard atomistic semantics, but something that, as in the case of "oval," can be identified by means of an implicit conception of what it is to be identical as something that is presently perceived as having some certain property. The only difference is that, now, the property is no more a monadic one.

Suppose I see two men, George and Mark, and that I perceive George to be taller than Mark. It is plain to see that any theory that assumes different perceptions as grounding sources for the references respectively of George and Mark's heights and the relation between them is not a good theory. For there is no different fact to be pointed at in perception. We can assume perhaps that relations merely supervenes on perceptibles, so that the only properties that are grounded in perception are George and Mark's heights, and then I derive that George is taller than Mark. But this is hardly what happens, since no inference is involved in such a recognition.

The case of relations, however, is not special. It is just an instance of a more general case. My perception of the oval shape of an object is never

¹¹ *Ibid.*, pg. 267.

detached from a perception of, say, something that is coloured. Nonetheless, in my reference to something oval, I know that I am abstracting from its colour. This can be explained only if, within our perceptual interactions, we assume some general taxonomic form, which enhances our referential competence, together with all the rational courses of actions and interactions that are part of our rationality.

Beyond Rationality Nothing of what I said so long about the central role of identity and reference should sound utterly new, for it is merely a re-proposal of the theory of previous chapters. But dealing with Peacocke's theory is important also because it provides us with clarifications to my past argumentation concerning self-interpretation.

To establish the meaning of sentences concerning states of mind, or mental actions, Peacocke assumes a privileged access provided by first-person assertions and attributions of such states of mind and mental actions, and then the recognition that other people are selves in the same sense in which I am a self: i.e., they have the same privileged access to their own states of mind and mental actions that I have to my own. Recognition of these facts (my access to my states of mind, the fact that this access is due to the possibility of using the first person, and, therefore, that each person capable of using the first person like I do is a subject in the very sense I am one) are enough to explain the interpretation of all talk about mental states and mental actions.¹²

To understand to what extent Peacocke's theory is in agreement with my own, and to what extent, instead, is in conflict with it, let us begin with what he describes as a "thin' account of what it is to have the conception of a range of subjects that includes oneself:"

A range of subjects is conceived of as a range of things of the same kind as me, standing in the same kinds of relations to the world, events, objects, and actions as are required for me to exist.¹³

The reason why such an account is called thin is, I think, that it is neutral as to the priority of conceiving of oneself as a subject with respect to conceiving of other people as subjects.

If, for one reason or another, you are sceptical of the existence of an intermediate level of minimal subject-involving thought, you could still consistently accept this thin account. You would just be committed to a more holistic understanding of the specification of the content of the thin account, an understanding according to which the first person and the conception of a range of subjects are explained simultaneously.¹⁴

¹² *Ibid.*, chpts. 3 and 5.

¹³ *Ibid.*, pg. 183.

¹⁴ *Ibid.*, pg. 184.

This is where my theory, as expressed in § 4.3, and Peacocke's agree. I also agree that when others talk about their sensations, or mental states or actions (like taking decisions, or struggling to believe) we understand them by projecting our experience of thinking beings and rational subjects upon them. There is a difference in this projection that should be noticed, though, for I think that it is the whole of the rational system that is attributed to other people, and that is how we can make sense also of sensations or mental states or mental actions that we have never been quite into, as long as we can see how it is that it is rational for our interlocutor (or for the rational being that we figured our interlocutor to be) to be into them or to entertain them.

The major point of conflict between my projective theory and Peacocke's, though, is in the role of the first person. Even if Peacocke's thin account leaves open the possibility of a simultaneous identification of oneself and other people as subjects, the access to the interpretation of talk about mental states and actions is always mediated, in his theory, by the self-attribution (via the first person) of such mental states and actions. Vice versa, such self-attribution is an immediate and privileged access to the meaning of expressions dealing with mental states and actions.

In § 4.3 I proposed instead that self-attribution is interpretable because of our sharing of the same cognitive system with our rational peers (our interlocutors), and our talk of mental states and actions is interpreted in no different way than that in which we interpret other people's talking of the same subject; namely, we know when our own utterances mean such and such mental state or action because we manage to identify the position of such utterances in our own reconstructed rational way of living.

It is important to notice that I do not want to deal with the problem, which seems instead rather central in Peacocke's theory, of the privileged access one has to one's own mental states and actions. What I am interested in is the mere verbalization of such access, though I suspect that such verbalization has an essential role for what concerns general access.

The point is to explain the reference of the first person, and possibly to explain it without circularity. For if we appeal to the privileged access that we have, as subjects, to our mental states and actions, i.e., to that peculiar feeling of acknowledging things from within, we must be assured, first, that the peculiar role of the first person in our languages has nothing to do with that. This is a difficult task, because it is not at all clear neither whether this is in fact the case, nor how we could realize that, if it were.

My opinion is that here we are bordering a great problem about the foundation of any theory of understanding whatsoever, especially those which appeal to some primitive notion of holistic rationality, such as mine. For I suspect that the first person (i.e., utterances of sentences involving the use of the first person) is constitutive of our knowledge of what it is to be a subject; in a sense, then, it is not appropriate to ask whether such knowledge is prior to the first person or vice versa. They are, so to speak, two aspects of the same

global question, that concerning our own rational life as subjects.

The problem is, then, about the foundation of such rational life, and of our reliance on principles such as the principle of charity or the principle of rationality as guides for our interpretations. It is difficult to understand in what sense such principles can act as regulative principles, as axioms, and yet be perceived, as per the Q/D thesis, as in some sense revisable. My answer is that they are indeed not revisable, and in this sense they are deemed as utterly *a priori*, but this is so only within our human and rational practice.

As human beings, we cannot help regarding our own lives as regulated by a rational principle, and this means interpreting our doings as informed by such a rational principle. At the same time, as if outside proper philosophy, we cannot help wondering about those beings that we see as excluded by such a rational principle, whether there is a quasi-rational principle for them, or even whether they possess their own special rationality, which is not only unknown to us, but utterly alien to the kind of rationality that regulates our actions and interpretations. As philosophers of mind, I think, we must resist this way of speaking, since there is no rationality outside rationality, and there is no coherence outside rational coherence, i.e., outside the coherence that we can understand as such.

I think that it would be inappropriate to conclude, then, that there is nothing at all outside our rational experience. It is not that there is something, admittedly, but there is not nothing, too. Outside our rational experience lies the realm of silence.

I admit that this conclusion may sound unsatisfactory to many, and rather mystical to most. This is perhaps a problem of holistic theories of interpretation: having to deal with the rationality that lies at the core of human nature as a whole and shows itself into action, holism cannot help surrounding rationality with mystery concerning its roots. I do not think that such a conclusion about rationality is avoidable by adopting a different semantic theory rather than holism, nor I feel forced to state that such a conclusion should be avoided at all. After all, it may very well be that the only answer to the question concerning the reason why we are the kind of beings that we are is, “by chance.” I do not feel compelled to find a better answer now; I only wanted to note that such a universal problem is quite soon encountered digging for the foundation of a holistic theory of understanding, while, perhaps, atomism and molecularism are better at ignoring the matter.

6.1.2 The Problems with Holistic Semantics

For what concerns the properly semantic part of the holistic theory of meaning, we still may have doubts regarding the feasibility of a theory trying to link the meaning of any single utterance to the architecture of the whole linguistic system. First of all, we might infer from the distinction between a theory of interpretation and a theory of meaning in a narrow sense, that the two

faculties are somehow distinct themselves.

If this is true, we may ask what is the role of the faculty of meaning and the conceptual faculty (which are the two main faculties we individuated in our discussion of holistic semantics, at least according to one kind of theory), as opposed to the interpretive faculty. In fact, it seems to be clear that the only faculty which has an observable effect on our behavior and on our cognitive activity is the interpretive faculty. After all, assuming for the moment that the other faculties indeed exist, and that our theory of interpretation is generally correct, any activity taking place in our cognitive system (hence, also activities of the linguistic and conceptual faculties) is filtered to our consciousness by the interpretive faculty.

Now, there is a reason why we need to keep distinct the interpretive system and the semantic system, although they are all part of the rational cognitive system which is assumed as fundamental for the functioning of at least the interpretive process. For the interpretive system, as such, is not a productive system, in the sense that interpretation is not enough to explain the production of always new meaningful sentences. It merely guarantees that for any novel sentence we should encounter, we are capable of finding an interpretation adequate to the circumstances in which it is uttered.

When we start to think about the semantic system (either with its division into a properly linguistic and a conceptual system, or as a complex global system which unifies the two aspects) as the productive side of our linguistic mastery, we cannot help asking whether there is a more reasonable way to deem the productive and the interpretive side of the matter as merely two ways of thinking about the same faculty or system, rather than two different systems working together. After all, there is an important sense in which our interpretation strictly depends on our capacity of producing the utterance with the meaning we are attributing to our interlocutor.

While this is true (at least according to the theory detailed in the past two chapters), we cannot ignore that the faculties of producing and interpreting meaningful utterances are sometimes detached. Some aphasic subjects, for instance, show a double dissociation between them, some being able to understand speech but not to produce it, while others are conversely able to produce correct speech, though incapable of understanding it. We can have the same kind of evidence concerning the distinction between a linguistic (mostly syntactic) system and a conceptual one; and of course we already quoted such evidence when dealing with the distinction between a referential and an inferential competence of the conceptual system.¹⁵

We still may wonder what all these empirical evidence is supposed to prove or show. My answer is that, outside a complex theory, it does not show or prove much. Only a theory about the architecture of our cognitive system can guide us towards an interpretation of empirical evidence. The other way

¹⁵Cf. § 5.2.4.

round is, in my opinion, simply wrong. While it is correct to expect different theories from different areas of research to dialogue, we cannot expect one discipline to dictate another one's agenda.

In this sense, such evidence is not at all conclusive for our attempt to reconstruct the human cognitive architecture. In fact, we are still faced with the dilemma we described in the previous chapter.¹⁶ We may think that these cases of double dissociation and the fact that we are able to distinguish different tasks developed by the cognitive system are hints of the existence of different subsystems – which, of course, would be interconnected and interdependent as for the developing of those tasks. Otherwise, we may think that there is only a cognitive system, and it is illegitimate to distinguish subsystems in it: we are only entitled to think of different tasks, but it is not the case that there is a biunivocal correspondence between tasks and subsystems. Perhaps even a third possibility is open, viz., that there are interconnected subsystems, and that, under certain circumstances, one can act as a vicarious for another one when the latter cannot execute its task.

6.2 Interactions

Towards the end of the previous section, I have introduced the argument of the interaction between holistic semantics on one side and other disciplines on the other side. We have already seen that there may be interesting connections between holistic theories of meaning and neurology or psychology. I think that it is even more interesting to see the place of holistic semantics with respect to other disciplines in the philosophical area (such as epistemology and theory of mind), and this will be my starting point. After that I will deal again with the connections between semantic holism and modern linguistics.

6.2.1 Philosophical Issues

For what concerns the relationship and the influence that a holistic theory of meaning has or may have with respect to other philosophical issues, Fodor and Lepore claim that there are reasons not only to believe that holism has not been proved true, but also to worry about the eventuality that such a proof should be given:

It appears that holism, taken seriously, entails incommensurability of theories, the impossibility of linguistic diachrony, and the impossibility of any generalization concerning the intentional states of people (thus the impossibility of psychology, economics, and so forth). According to Fodor and Lepore, these are all consequences of holism in their sense. [...] Anyway, it appears that the dreadful

¹⁶§ 5.1.

consequences that Fodor and Lepore draw [from their definition] could equally well be derived from holism in the traditional sense.¹⁷

All these problems roots in the same ground:

[...] meaning holism would require that if any one sentence in your theory occurs in my theory, then practically all the sentences that occur in your theory must occur in my theory. An similarly, mutatis mutandis, if “theory” is replaced with “language.” If holism is true, then I can’t understand any of your language unless I can understand practically all of it.¹⁸

Let us go back to § 4.1.2, and see what was the problem with Argument A, in Fodor and Lepore’s opinion. They claim that their counterargument shows that the same belief *P* can be shared by me and you, even if it is related to different other beliefs in my mind and in yours (say, by the beliefs in the set *A* for me and by the beliefs in the set *B* for you). But if holism claims that to share any one belief one has to share also lots of other beliefs, then there is a contradiction between holism and Fodor and Lepore’s counterargument. I agree that there is a contradiction, but I do not see why we should decide that it is holism that should be abandoned because of the contradiction. We could as well decide that the hypothesis that we were able to identify a belief independently of the context furnished by other beliefs is untenable. After all, Fodor and Lepore have not given us an alternative method of belief individuation. They were not supposed to, of course, because the aim of their work is not to furnish an alternative theory to holism. However, a counterargument based on premises that are not shared from the thesis it is supposed to attack is doomed to miss the point.

Fodor and Lepore think that their counterargument works, because they claim to have independently opposed to the holistic doctrine that the individuation of semantic entities depends on the whole semantic system. They think, in fact, that such a doctrine is directly responsible for untenable consequences, which are summarized by Marconi in the quote above. I would like, therefore, to discuss such consequences of holism, and its defense against the charges moved by Fodor and Lepore, in order to clarify the connections between semantic holism on one hand and epistemology and theory of concepts on the other.

Here is how I will procede. I will first examine the “dreadful consequences” of holism more in detail, to see where their roots really lie, and whether holism can deal with them. Then I will see what holism is claimed by Fodor and Lepore to have responded to such allegations (allegedly with no success). In

¹⁷Marconi (1997), pgs. 48 f.

¹⁸Fodor and Lepore (1992), pg. 9.

examining this defense strategy, we will see whether some method of holistic individuation to dismantle Fodor and Lepore's case against Argument A appears.

Two Faces of Evil

The two consequences I intend to examine in this section are incommensurability, for what concerns the connections between semantic holism and epistemology, and the impossibility of intentional generalization, concerning the connections between semantic holism and conceptual theory. Of course, these are not the only connections among these disciplines: some others have appeared in the previous arguments detailing semantic holism. These were left aside, however, and this is the reason I deal with them here.

It will appear that I have different intents in the following argument. First of all, I will see whether Fodor and Lepore's derivations of the dreadful consequences are sound; in other words, I intend to see whether they are consequences at all. Second, I will be concerned with their alleged dreadfulness: even if semantic holism led to such consequences, are Fodor and Lepore correct in thinking that this is enough reason to doubt of the sustainability of semantic holism? It could turn out either that there is no reason to be afraid of these consequences as such, or that, dreadful as they are per se, semantic holism effectively defuses their most dreadful aspects.

Incommensurability In deriving incommensurability (which is a theory concerning philosophy of science) from semantic holism (which is a theory concerning philosophy of language) one attempts to make a leap from languages to theories. Such a leap is not always a safe one. Let us see, then, how Fodor and Lepore jump. First, the linguistic side of the chasm:

Consider the property R that a linguistic expression has iff it refers to the same thing that some expression in English does. [...] Question: Is the property R holistic? Could languages that overlap only slightly share any of their "ontological commitments?"

Here's one reason why this question matters. Suppose that ontological commitments are holistic, so that two languages can share any of their ontology only if they share quite a lot of it. It might turn out, for example, that no language could have an expression that refers to what the English expression "the pen of my aunt" refers to unless it also has expressions that refer to, as it might be, Chicago, the cat's being on the mat, the last game of the 1927 World Series, the day after they built the Statue of Liberty, the last of the Mohicans, *The Last of the Mohicans*, and so forth.¹⁹

¹⁹ *Ibid.*, pg. 11.

First of all, let me state that it is not at all clear what here is intended with “ontological commitments.” Languages, as such, have no ontological commitments at all, or it is at least disputable that they do. It seems to me that it is preferable to say that speakers can have their own favourite ontological theory according to which they interpret their own language and other people’s language. And the same language, so it seems, can be interpreted according to different ontological theories.²⁰ And this seems to be the reason of all the ontological disputes across the millennia: that philosophers could not assess the question simply by looking at the language. I will put all the matter in a dubitative form, but I would like to note that there is some danger on Fodor and Lepore’s leap: the pole they meant to use may be less solid than it appears.

As a matter of fact, I do not find anything dreadful in Fodor and Lepore’s hypothesis.²¹ I would indeed be puzzled if I was told that someone could refer to a pen but not to a city, to a book (either the physical object or the story wherein told), to a particular individual, to a time, to a state of affairs, or to an event. I would be puzzled because I would have thought that, these things being part of what is or could be relevant for any human being to be talking about, any human being had means to refer to them. All this may simply boil down to saying that my intuitions are different from Fodor and Lepore’s. It may be that some empirical inquiry about human beings and human languages can prove me or them right by showing whether there can exist languages lacking the expressive means to refer to any of the above mentioned entities. I doubt that such empirical evidence can be found, though, unless we prove that there is a univocal way to attribute ontological commitments to other people.

This having been said, let us see the final phase of the leap:

It raises the stakes [...] that the same considerations would apply if we asked about the semantic property R^* . An expression has

²⁰Cf. Davidson (2001), pgs. 188 f.:

Suppose that in my office of Minister of Scientific Language I want the new man to stop using words that refer, say, to emotions, feelings, thoughts, and intentions, and to talk instead of the physiological states and happenings that are assumed to be more or less identical with the mental ruff and raff. How do I tell whether my advice has been heeded if the new man speaks a new language? For all I know, the shiny new phrases, though stolen from the old language in which they refer to physiological stirrings, may in his mouth play the role of the messy old mental concepts.

The key phrase is: for all I know. What is clear is that retention of some or all of the old vocabulary in itself provides no basis for judging the new scheme to be the same as, or different from, the old.

²¹Nor do they say they do, of course, apart from finding it “counter-intuitive.” Cf. Fodor and Lepore (1992), pg. 11.

R^* iff it refers to something or other that currently accepted astronomical theories refer to. Suppose that R^* is anatomic, hence holistic on the assumption that anatomism implies holism. Then it might turn out that no theory could refer to (for example) stars unless it could also refer to (as it might be) planets, nebulae, black holes, the center of the galaxy, the speed of propagation of light, and the location of the nearest quasar. It would follow that Greek astronomy (hence, Greek astronomers) couldn't ever refer to stars. [...] In fact, strictly speaking, it would follow that the Greeks didn't *have* any views about *stars*; we can't, in the vocabulary of contemporary astronomy, say what, if anything, Greek astronomy was about. A fortiori, it makes no sense to speak of an empirically motivated choice between Greek astronomy and ours; whereof you cannot speak, thereof you must be silent.²²

Again, my intuitions part company with Fodor and Lepore's, for I do not see any inconvenience in such a proposal. I am not alone in this, but I will not try to prove my party right in this occasion. Fodor and Lepore say that if we think that we can bite this bullet, we must also accept that science is not progressive, and that, therefore, "the standard argument for Scientific Realism goes down the drain."²³ I can be fine with all this, although I fail to see the cogency of the inference.

On one side, in fact, note that some comparison is still permitted even if we deal with incommensurable theories. Thomas Kuhn was certainly one of the keenest philosophers on incommensurability. He was, indeed, one of the first to have introduced the term in the epistemological debate. Yet he did not think that incommensurable meant that no comparison at all was possible:

Remember briefly where the term 'incommensurability' came from. The hypotenuse of an isosceles right triangle is incommensurable with its side or the circumference of a circle with its radius in the sense that there is no unit of length contained without residue an integral number of times in each member of the pair. There is thus no common measure. But lack of a common measure does not make comparison impossible. On the contrary, incommensurable magnitudes can be compared to any required degree of approximation. [...]

Applied to the conceptual vocabulary deployed in and around scientific theory, the term 'incommensurability' functions metaphorically. The phrase 'no common measure' becomes 'no common language.' The claim that two theories are incommensurable is then the claim that there is no language, neutral or otherwise,

²² *Ibid.*, pgs. 11 f., emphasis in text.

²³ *Ibid.*, pg. 12.

into which both theories, conceived as sets of sentences, can be translated without residue or loss. No more in its metaphorical than its literal form does incommensurability imply incomparability, and for much the same reason. Most of the terms common to the two theories function the same way in both; their meanings, whatever those may be, are preserved; their translation is simply homophonic. Only for a small subgroup of (usually interdefined) terms and for sentences containing them do problems of translatability arise.²⁴

This subgroup of technical terms can very well be considered responsible for the better predictions that enables us to say that our theory is better than the Greeks'. Notice how this result is achieved by means of a holistic theory of meaning, and not against it: if we manage to put in a (homophonic) relation terms of both theories is because they serve the same function in both theories. Also, note that a scientific theory is never independent from the human beings that made it up, and that it is also thanks to what we know about Greek culture and Greek language that we can say that what they were trying to devise was an astronomical theory.²⁵

As for Scientific Realism, the reason it goes down the drain may not be the fact that we cannot account for progressive science any more. Maybe the problem is that we do not really know what to do with a concept of real world as opposed to the scientific picture which is supposed to be its approximation. If scientific progress really was "*the* argument for Scientific Realism,"²⁶ maybe Scientific Realism was condemned anyway: indeed, there seems to be quite a lot of more independently well-equipped epistemological competitors that could account for progressive science. And some other competitors would also gladly reject the picture of a progressive science altogether.

In fact, Fodor and Lepore's leap from philosophy of language to philosophy of science was not only dangerous and, perhaps, even ill-supported by their mention of reference. It was also not at all necessary. Epistemological theories about the holism of scientific theories are independent from holistic theories about language.

Imagine we assume a rather classical atomistic point of view about meaning: the meaning of a sentence is the fact that verify it. Suppose further that

²⁴Kuhn (2000), pg. 35 f.

²⁵Cf. Davidson (2001), pg. 168:

It isn't that any one false belief necessarily destroys our ability to identify further beliefs, but that the intelligibility of such identifications must depend on a background of largely unmentioned and unquestioned true beliefs.

Note that this would be, in Fodor and Lepore's terminology, a shift from talking of content identity to content similarity. I will come back on this issue further on.

²⁶Fodor and Lepore (1992), pg. 12, emphasis in text.

we try to be more specific about that, and that we decide that what counts as verifying has something to do with our perceptions. We can think that it is a task for a scientific theory to show how all of our sentences are connected to our perception. This was a strong positivist claim, of course, but it was not alien to Quine, either.²⁷ Now suppose that, as it may turn out to be the case, whether a sentence is verified by a fact or not depends on the whole theory. This was of course Quine's opinion; remember, on this regard, that Fodor and Lepore mention their doubt about Quine being a semantic holist.

All my point is to show that, as Fodor and Lepore seem to acknowledge elsewhere, it is not at all necessary to be a holist about meaning to be a holist about scientific theories, with all the consequences of that kind of holism. Scientific holism can be seen as one of the natural consequences of the positivist program facing its extreme consequences.²⁸ On the other hand, as we have seen, semantic holists can have some good argument to block the radical and perhaps unwanted consequences of scientific holism.

The Impossibility of Intentional Generalization Another argument is built by Fodor and Lepore starting with translation instead of reference. Suppose that we assume that the property of being a translation of one sentence of English or the other is holistic. It will follow that if a person's language cannot translate lots of present English sentences, it also cannot translate, say, "My aunt's pen is on the table." The problem is, of course, that lots of names and lots of concepts happen to have been introduced during the centuries; so, e.g., Shakespeare, or Chaucer, or Lincoln could not translate "My aunt's pen is on the table," because they lack a translation for, say, "Water is XYZ" or "The proton is made of quarks."²⁹ This is why Marconi talked of "impossibility of linguistic diachrony" as a consequence of holism.

This part of Fodor and Lepore's argument seems very close to the first part of the argument considered in the last section, and I think that the comments I gave there still apply here. In a nutshell: I think that languages should be thought of as permitting to acquire new vocabulary with only slight adjustments to the already acquired vocabulary. The mechanism that should permit such slight adjustments has already been the subject of previous discussion.³⁰

The argument, however, is not finished yet. It carries on, claiming that, if an issue can be made about Shakespeare's language, it can be made also about his thoughts. For, consider the property of entertaining one or the other belief that I entertain. If this is holistic, then Shakespeare could not believe that the Thames is a river because he could not believe that Obama won the elections. The problem is that we assume that human actions depend on

²⁷Cf. Quine (1953), pg. 42.

²⁸This is why Rosaria Egidi can consider Hanson and Feyerabend, who mostly shared Kuhn's views, (dissident) neopositivists. Cf. Egidi (1979), p. 8.

²⁹Fodor and Lepore (1992), pg. 13.

³⁰See § 4.2.2 above.

what people believe and on what they want, and if there is no way of making generalizations about beliefs and desires, then there is no chance of having a predictive theory about human behaviour.³¹

This is because thoughts can fall under generalizations only if they are identified by means of their content. If content, in its turn, is holistically identified, then we cannot equate content of thoughts across minds, because every person has very different thoughts that others do not entertain.³² This, in turn, entails the third consequence dreaded by Marconi: the impossibility of scientific knowledge about intentionalities and, therefore, the impossibility of psychology, sociology, economics, and so on.³³

To this we could answer that holistic individuation (or, actually, quasi-individuation) of intentional content can be achieved by a trial and error procedure, which is behavioristic inasmuch as it takes behavioral outputs into consideration, but which is not entirely behavioristic inasmuch as it gets a feedback by confronting its partial results with the success of the linguistic interaction. In other words, we can begin postulating some thought to be present in our interlocutor's mind, then check if we pictured the right thought keeping trace of it in our further conversation. This seems the way that people come to the conclusion that they have the same thoughts as their interlocutors.

I have to note two things about this line of answer. First, to answer to Fodor and Lepore's worries, such a procedure gives enough way to generalizations. Of course there is the possibility of committing mistakes, and taking different thoughts as the same or vice versa. But the possibility of being mistaken is deeply rooted in every human activity – even in hard sciences –, so I would not be worried about that.³⁴ On the other hand, it could be objected that we are giving no identification criteria, but at best a procedure to decide whether two thoughts have more or less the same content. That is, we are no more talking of sameness of content, but of likeness of content, and here we could have other problems.

Content Similarity

In both the situations mentioned in the last subsection, we have seen that the holist would appeal to some criteria of content similarity to avoid or mitigate the consequences of being a holist. What is Fodor and Lepore's answer to

³¹Fodor and Lepore (1992), pgs. 13 f.

³²*Ibid.*, pg. 14.

³³*Ibid.*, pg. 15.

³⁴There is also another way in which so-called human science are similar to hard sciences. Fodor and Lepore (*ibid.*, pg. 17) seem to have problems with assumptions that allow human science to be “not negotiable” in the face of, e.g., biology or physics. But this is not a special status that, accepting a holistic point of view, intentional explanation would have, in contrast with physics and biology. For one might hold, for instance, chemical explanations to be “non negotiable” however physics turn out to be. Indeed, this is what happened in the 18th and 19th centuries.

that? They say that you cannot have any opinion about content similarity unless you know how to identify beliefs in the first place – and, maybe, how to count them. Here is an example where these problems come out:

No doubt, one does know (sort of) what it is like to more or less believe the same things as the President does; it's to share *many of the President's beliefs*. For example, the President believes P, Q, R, and S, and I believe P, Q, and R; so my beliefs are similar to his. An alternative, compatible reading is: the President believes P and Q *very strongly* and I believe them equally strongly or almost as strongly, so again my beliefs are similar to his.³⁵

If this were the right way to construe “more or less believe the same things as the President does,” Fodor and Lepore would be right in saying that such a notion depends on the notion of identity. However, we can think that, once again, they are misled by their intuitions. The naive picture that they draw needs to be questioned more than they do, before we can say anything about it. The question I would like to ask resembles the question that Fodor and Lepore seem to ask, viz., How do we know that we believe more or less the same believes that the President does?

I would read that question as asking: In what way we come to have that knowledge? If I were asked, I would answer that I have listened to, or read about, the President's speeches, and this is how I know what he believes. A theory of interpretation is involved here. Maybe also a practical theory, a theory that takes actions into consideration.

Fodor and Lepore could object that we cannot have any such theory without generalizations about intentional content, and that such generalizations rely on content identification. Actually, we have seen before that this might not be necessary. Someone might object that, when we showed that, we did so by substituting content identity with content similarity; thus, if now we are trying to support content similarity with the same move, we are begging the question.

I do not think that it is so: previously, we have announced the possibility of a structured scientific methodology. Here we are assuming an explanation of our intuitions. Both scientific methodology and intuitions rely on an explanation of content similarity. This, in its turn, is grounded on a theory of interpretation, which, having the principle of rationality as one of its constitutive principles, explains very well the possibility of intentional generalization. And here there is no begging the question, for the principle of rationality is allegedly primitive.

What I would like to ask is whether my account is more accurate than Fodor and Lepore's with respect to being an account of what it is like to more or less believe the same things as the President does. Remember, though,

³⁵ *Ibid.*, pg. 18, emphasis in text.

that Fodor and Lepore have just a negative task, viz., to show that holistic positions have not yet been plausibly defended. What I wanted to do, instead, is to show that there is some plausibility in the picture I have given in the present section – which is, of course, backed up by the theory examined in § 4.2. If it is so, then the counterargument proposed by Fodor and Lepore and discussed in § 4.1.2 to what they call Argument A in favour of semantic holism is proven definitely wrong, because it is apparent that it was based on assumptions that holists have no reason to share.

In this section, I did not want to claim that semantic holism is a key to the solving of the several philosophical problems ranging from ontological question to epistemological ones and to those concerning the theory of mind. I rather think that the present discourse can show that semantic holism can give us a peculiar perspective on these issues, placing us in a route to the solutions. At the same time, I also acknowledge some degree of independence of these issues from semantic ones, in the sense that our commitment to semantic holism (or to other kinds of semantic theories) does not equally commit us to the kinds of solutions described in the present section (or, respectively, to necessarily different solutions).

6.2.2 Linguistics

We have seen throughout the previous chapter that semantic holism is a great premise to the whole construction of modern linguistic theories. Therefore, semantic holism is entailed by linguistics as it is practiced nowadays. Once again, it is important to state clearly what the relationship between semantic holism and modern linguistics (or between any philosophical theory and any scientific theory, for that matter) consists in.³⁶ When I say that the modern practice of linguistics presupposes certain philosophical assumptions such as semantic holism, I do not want to imply that the success of modern linguistics in explaining the functioning of the human faculty of language, or the confirmations that may come to modern linguistics from experimental evidence, are to be deemed proofs of the correctness of holistic theories of semantics.

As I said before,³⁷ the interpretation of experimental evidence depends strictly on the theory (therefore, on the theorists) that attempts the interpretation. As for the relationship between a theory and its assumptions, it

³⁶It is perhaps once again the case to clarify also that the distinction between what is to be deemed a scientific theory and what is not is not so sharp as we would think. To debate this issue here would lead me outside the boundaries of the present work, so I will just note that there is a problem, without attempting to solve it. Anyway, the solution is largely irrelevant to the point I want to make, because the relationship between philosophy and science, as I want to describe it here, merely depends on what we judge to be philosophy or science at a given time, independently on the reasons and justifications that we might have in doing so.

³⁷Cf. § 6.1.2.

is basic logic that teaches us how the conclusion might be true even if the premises are false. Of course, if we manage to show that the conclusion is not derivable from other premises, then we have stepped further towards the ascertainment of the premises; in this case, of the philosophical assumptions which ground modern linguistics. It must be noticed, though, that we can never be assured of the certainty of the conclusion itself (i.e., of the scientific theory), and this for the very nature of scientific enterprise. So, at best, there is a mutual relationship of reinforcement between semantic holism and modern linguistics, but it is not that they stand or fall together, nor that one of the two is necessary or sufficient for the truth of the other.

This relationship and reinforcement can be best understood, in my opinion, if we focus on the heuristic value that semantic holism may have with regards to some scientific enterprises. Far from having the deleterious effects which Fodor and Lepore would attribute to it, semantic holism can prove a useful tool for some scientific disciplines – linguistics being only one of them – inasmuch as it can show where theoretical problems are, as well as which routes are more likely to provide a solution, given the assumptions of the discipline. My intention in what follows is to give some example of this heuristic power of semantic holism with respect to contemporary linguistics.

Relevance Theory

We have seen some of the fields of linguistics in which some holistic theory is common background. Syntactic theories, for instance, rest by and large on holistic assumptions. So does any generativistic theory: since the generativistic approach originated first as a syntactic theory, we can expect that it has retained some flavour from that.

Also pragmatics can be favoured by a holistic approach. We have seen, in § 3.3.2, that some pragmaticists would be glad to have some theoretical background to help them individuate what is relevant to the practical uses of language. The problems that we have found in dealing with a molecularistic theory, which we might have expected to ground such an individuation, can lead the pragmaticists to look somewhere else for the background they might need. Semantic holism, with its emphasis on what cognitivists call ‘theory of mind’ (i.e., the attribution of mental states to other people as a crucial step in explaining their behaviour) and on the importance of focusing on the whole experience of oneself as a rational being, seems pretty much in agreement at least with some of the most credited theories in pragmatics, such as Relevance Theory.

According to this, human beings are compelled, in their uses of language, to regard all it is said as being said by the speaker for the purpose of being relevant for the hearer. With this axiomatic ground in mind (which clearly resembles the principle of charity), one calls on everything known about oneself and the interlocutor to make sense of the utterances that have been spoken.

It is a trial and error procedure, and its final point is the establishing of the amount of modification to individual knowledge that the processing and interpretation of an utterance causes to the hearer. If such modification is not worth the interpretive effort, then there is no relevance to be found in the utterance. Even if this happens sometimes (for not every chunk of speech that we happen to hear is relevant to us), the hearer cannot doubt the speaker's intention of being relevant, unless the doubt is casted onto the speaker's rationality altogether. Rather, the hearer would look for some kind of mistake from the speaker – in the felicity of the utterance as an expression of the speaker's intention, or in the speaker's grasping of the interactive situation, etc.³⁸

Sketchy as my presentation might be, it is still easy to see how this theory appeals to the same background theory as detailed in § 4.2.7, i.e., on the attribution to the speaker of the same cognitive system of the hearer as a whole. Relevance, in fact, can be determined only against a background of shared belief, and the attribution of shared belief is possible, according to the account that I have been presenting, only within a holistic theory of interpretation. In this case, then, a theory about pragmatics (Relevance Theory) and a theory of philosophy of language (semantic holism) find an agreement, perhaps through rather independent routes. Still, this agreement shows that it is possible that one theory points problems or solutions to the other, which would perhaps pass unseen, were the theories left alone in their respective fields.

Acquisition

We have already seen some issues about the problem of linguistic acquisition that can be solved by appealing to some holistic theory of language.³⁹ We have passed under silence, though, what might seem the major obstacle posed by the phenomenon of linguistic acquisition to holism.

Here is the objection, as it might be stated. It is easy for common sense to imagine how, one step after the other, children expand their languages into adult ones, adding one piece at a time to a pre-existing structure. But, if the child's language is part of the adult's one, then holism must be false. The intuitive picture is not an obstacle because it pictures any chunk of language as independent from the others. After all, it is easy to imagine that adding something to a structure modifies the structure itself, creating the network of interdependencies that is needed by a holistic theory. The problem is, rather, that we have repeated many times that, under holism, we should regard the whole language as primitive, and the parts and elements of it as derivative; but the intuitive picture of acquisition clearly subverts the holistic image, as we conceived it.

It seems from independent arguments, though, that holism holds. Hence, we might be tempted to look for an explanation according to which, against

³⁸Cf. Sperber and Wilson (1986) for further details.

³⁹Cf. § 5.2.3.

appearances, the child's language is not part of the adult's, but it is a different linguistic system. Our problem is to explain what the connection between the two linguistic systems, the child's and the adult's, could be. We must find an alternative conception of acquisition to the gradual learning process picture.

Chomsky's idea is that in the case of language we have something analogous to what happens in the case of the immune system for what concerns selection of antibodies by antigens.⁴⁰ We might think that, in the presence of a particular antigen, the human body reacts producing the correspondent antibody. Actually, this is not the proper way of describing what happens. The body already has the antibodies it might need: the antigen just selects the correspondent one. In the same way, children are born with all grammars they might need, and their linguistic environment selects the right one. To put it in a better way: they are born with a universal grammar which shows them how to fix certain parameters starting from linguistic data they come in contact with to obtain the grammar which generates those linguistic data. One should not talk, therefore, of partial grammar: grammar is complete from the beginning. From the beginning a structure exists and interacts with other structures, with the environment, and with genetic predisposition, and thus it performs a certain function and reshapes itself into another structure, and so on, until it reaches a stable stage – i.e., adult grammar.⁴¹

Saying that a child's grammar is part of an adult's grammar is as senseless as saying that its brain is part of the adult's. It is a preceding stage of it, but it does not exist because of what it will turn into: in fact, it is not at all determined which direction its development will take. This holds for brain as much as for language, and for the whole human body.

Now I turn once again to heuristics. We saw that to assume a holistic point of view on language gives us a theory of acquisition, which can and has to be tested against empirical evidence. Contrasting theories offer less room for empirical test.

We already mentioned the problem of vagueness of molecularism, vagueness which makes all tests non-decisive and perhaps uninterpretable. There is another problem with molecularism about acquisition: it does not explain how molecules are acquired. In fact, if they are acquired one piece at a time, it is, actually, a form of atomism. If they are acquired each as a whole, again I see no remarkable difference from a holistic theory that should make us willing to prefer a molecularistic one. Both holism and molecularism, in fact, must explain acquisition of language in terms of acquisition of rules. A molecular-

⁴⁰Chomsky (1980) pp. 136 ff.

⁴¹Cf. *ibid.* p. 134:

Without attempting to inquire into too many subtleties or to settle the question, I would like to suggest that in certain fundamental respects we do not really learn a language; rather, grammar grows in the mind.

istic theory, moreover, should also contemplate metarules to choose among molecules: acquisition of these metarules must be explained, and this reduces molecularism's appeal.

An atomistic theory of acquisition has its difficulties too: it has to dismantle a sorites. We already proved that a single sentence means nothing in isolation. We can easily maintain that the same holds for two sentences in isolation. If acquisition were gradual, some sentence at a time, one should explain when and how meaning comes in the picture.

These problems might not be insurmountable. With time, labour, talent and luck someone could mend the faults. I think that, meanwhile, we should admit that the best theory we have is, for heuristic reasons, semantic holism.

Language Evolution

Another field of linguistic research in which semantic holism proves a powerful heuristic tool is the investigation of language evolution. If the faculty of language is, as it is thought by most prominent contemporary linguists, an innate characteristic of human beings, and if this translates in it being a biological feature of human beings, it is appropriate to ask about the evolutionary origin of such faculty.

Of course, it is not at all guaranteed that the faculty of language is biological in the first place. What we know from the discourse of the previous two chapters is that it is essential, for our theory of meaning and interpretation, to assume that the faculty of language is a universal characteristic of human beings. What this universality might consist in is completely another question. Anyway, since the study of such faculty is dealt by linguistics, and since modern linguistics, after Chomsky, poses itself as a branch of psychology, we might be tempted of finding a grounding to such a universal characteristic which is not metaphysical, but immanent and contingent in nature. Otherwise, we are forced to establish some sort of metaphysical argument to explain why human beings are endowed of such a faculty, while other kinds of beings are not.

Such an argument would be, as I said, necessarily metaphysical in character. Of course, also the claim that the faculty of language is biological, and therefore we are entitled in searching for its evolutionary origin is a metaphysical claim; at least in the sense that the only proofs (or hints) for the truth of such claim can come from the scientific research which is prompted and grounded from the claim itself. I do not think that this situation is particular hard to swallow, nor that it is rare. On the contrary, I think that it happens quite often in scientific enterprises to ground on unproved claims entire fields of research which are supposed to show the truth of the claims themselves. As long as this methodology is accepted by the community of researchers and it produces results which are equally accepted by them, I do not think that we should frown at it.

The long and the short of this discourse is that nowadays the leap from considering some certain cognitive characteristic (viz., the faculty of language) a universal feature of some organism and considering it a biological feature of that organism is seen as legitimate, even if it is not warranted by what we know about the connections between biological and cognitive features. Within such a leap, semantic holism has an important role, being the theory which assess the faculty of language as a universal feature in the first place.

The problem of the evolutionary source of the faculty of language comes to the attention of the scientific community from reasons that are not completely dependent on semantic holism, and that I will not attempt to reconstruct here. What I would like to show, though, is that also for this scientific project semantic holism can constitute a valid heuristic theory, pointing at routes that might seem promising to explain how the faculty of language evolved.

I will not go into any detail here: suffice it to say that the Davidsonian theory of interpretation, according to which interpretation always entails attribution of intentions, can be seen as showing a strong connection between the faculty of language (or, at least, that part of the faculty of language that is responsible for interpretation) and that faculty that cognitivists call ‘theory of mind.’ A promising route for the evolutionary linguistics is therefore to investigate the possibility that the faculty of language primarily evolved from the theory of mind. Furthermore, the suggestion that all interpretation is possible because of the ability of producing utterances, and the proposal that all self-interpretation comes first from recognition and interpretation of others’ intentions also strongly point in the same direction. These two ideas, taken together, also show that it is perhaps appropriate to think of the faculty of language as unitary, and deem the interpretive side and the productive side merely as two interrelated tasks carried out by the same faculty. If this is correct, then we might also have the hint that the interpretative task is prior to the productive task, and this may originate a very particular kind of evolutionary theory of the faculty of language.

It is interesting to note that a theory along these very lines has indeed been proposed. In a nutshell, the faculty of language first originated from the need of interpreting very long chunks of incoherent oral noises. The need of interpreting them, and of projecting the interpretation on the proto-utterances gave to them arrangements and orderings that eventually led to syntax and morphology.⁴²

⁴²Cf. Aitchison (1998) and McDaniel (2005). Of course, this is too little space to outline all the details of such a complex theory as one explaining the evolution of the faculty of language. For a collection of seminal studies in the field, cf. at least Hurford et al. (1998) and Tallerman (2005), where the two above-mentioned papers can be found.

6.3 Final Remarks

In this work I proposed to analyze the notion and theory of semantic holism, and I attempted to do it by comparison with other theories which intend to explain the same class of phenomena as semantic holism. In sketching the three theories and trying to see what functions in them and what arguments still need some mending, an eye was always kept on the relation between philosophy of language (as instantiated respectively by semantic atomism, semantic molecularism, and semantic holism) and other philosophical or scientific disciplines. To each of the three theories I grant some heuristic power, more or less so depending on the discipline we are interested in drawing the connection with.

Thus we have seen that semantic atomism seems apt for grounding formal logic and formal semantics, and it both grounds and depends on certain ontological assumptions. Semantic molecularism and semantic holism are perhaps best understood as trying to offer more or less the same kind of answer to the semantic problem. It is, for both cases, an answer which does not rely on what there is, i.e., there is no need or little need for ontological assumptions. Rather, it is quite necessary, in order to ground such theories, to question epistemology, therefore both theories have strong interdependencies with this discipline. As for the scientific connections, both have some suggestion to give to and to take from psychology and perhaps neurology. Also, the strongest connection for both disciplines is with linguistics.

While I gladly admit that, especially for the grounding of psychology and pragmatics, semantic molecularism is the kind of theory that we might need to devise in the long run, presently it has flaws at its foundation that can be mended only by paying attention to the objections coming from semantic holism, and to the kind of solution that semantic holism permits. If I was asked to foretell what kind of theory our semantics will evolve into, therefore, I would guess it will be some kind of holism reproducing for practical reasons the desirable features of semantic molecularism – e.g., the possibility of restricting in a straightforward way (which is still to be found) the field of what is relevant for a given task of interpretation. Or, maybe, something completely different.

Bibliography

- Aitchison, J. (1998). On discontinuing the continuity-discontinuity debate. In Hurford et al. (1998).
- Bermúdez, J. L. (1998). *The Paradox of Self-Consciousness*. the MIT Press, Cambridge MA.
- Bilgrami, A. (1992). *Belief and Meaning*. Blackwell, Oxford UK.
- Borer, H. (2005a). *Structuring Sense*, volume I. In Name Only. Oxford University Press, Oxford UK.
- Borer, H. (2005b). *Structuring Sense*, volume II. The Normal Course of Events. Oxford University Press, Oxford UK.
- Cairns, P., Shillcock, R. C., Chater, N., and Levy, J. (1994). Modelling the acquisition of lexical segmentation. In *Proceedings of the Child Language Research Forum*, Chicago. CSLI, Stanford CA, University of Chicago Press.
- Celata, C. and Calderone, B. (2008). Restrizioni fonotattiche, patterns lessicali e recupero delle regolarità morfologiche. Evidenze computazionali e comportamentali. In Bertinetto, P. M., Bambini, V., and Ricci, I. e. a., editors, *Linguaggio e cervello/Semantica, Atti del XLII Convegno della Società di Linguistica Italiana*. Pisa, Scuola Normale Superiore, in press.
- Chomsky, N. (1969). Quine's empirical assumptions. In Davidson and Hintikka (1969).
- Chomsky, N. (1980). *Rules and Representations*. Columbia University Press.
- Chomsky, N. (1986). *Knowledge of Language. Its Nature, Origin, and Use*. Praeger Publishers, Westport CT.
- Chomsky, N. (2000). *New Horizons in the Study of Language and Mind*. Cambridge University Press.
- Cutler, A. and Butterfield, S. (1991). Word boundary cues in clear speech: A supplementary report. *Speech Communication*, 10(4):335–353.

- Cutler, A. and Butterfield, S. (1992). Rhythmic cues to speech segmentation: Evidence from juncture misperception. *Journal of Memory and Language*, 31(2):218–236.
- Cutler, A. and Norris, D. (1988). The role of strong syllables in segmentation for lexical access. *Journal of Experimental Psychology: Human Perception and Performance*, 14(1):113–121.
- Davidson, D. (1980). *Essays on Actions and Events*. Clarendon Press, Oxford UK.
- Davidson, D. (2001). *Inquiries into Truth and Interpretation*. Oxford University Press, Oxford UK, 2nd edition.
- Davidson, D. and Hintikka, J., editors (1969). *Words and Objections. Essays on the Work of W. V. Quine*. Reidel, Dordrecht.
- Davis, M. H. (2003). Connectionist modelling of lexical segmentation and vocabulary acquisition. In Quinlan, P. T., editor, *Connectionist Models of Development*. Psychology Press, Hove UK.
- Devitt, M. (2006). *Ignorance of Language*. Oxford University Press.
- Dummett, M. (1976). What is a theory of meaning? 2. In Evans, G. and McDowell, J., editors, *Truth and Meaning*. Clarendon Press, Oxford UK.
- Egidi, R. (1979). *Il linguaggio delle teorie scientifiche*. Guida, Napoli.
- Everett, D. (2005). Cultural constraints on grammar and cognition in Pirahã: Another look at the design features of human language. *Current Anthropology*, 46.
- Everett, D. (2007). Cultural constraints on grammar in Pirahã: A reply to Nevins, Pesetsky, and Rodrigues. <http://ling.auf.net/lingBuzz/000427>
- Everett, D. (2009). Pirahã culture and grammar: A response to some criticism. *Language*, 85(2).
- Fodor, J. and Lepore, E. (1992). *Holism. A Shopper's Guide*. Blackwell, Cambridge MA.
- Frege, G. (1953a). *Collected Papers on Mathematics, Logic, and Philosophy*. Blackwell, Oxford UK.
- Frege, G. (1953b). Die Grundlagen der Arithmetik. Eine Logisch-mathematische Untersuchung über den Begriff der Zahl. In Frege (1953a).
- Frege, G. (1953c). Function und Begriff. In Frege (1953a).

- Frege, G. (1953d). Über Begriff und Gegenstand. In Frege (1953a).
- Frege, G. (1953e). Über Sinn und Bedeutung. In Frege (1953a).
- Goldsmith, J. (2007). Towards a new empiricism. *Recherches Linguistiques à Vincennes*, 36.
- Gordon, P. (2004). Numerical cognition without words: Evidence from Amazonia. *Science*, 306.
- Gow Jr., D. W. and Gordon, P. C. (1995). Lexical and prelexical influences on word segmentation: Evidence from priming. *Journal of Experimental Psychology: Human Perception and Performance*, 21(2):344–359.
- Gow Jr., D. W., Melvold, J., and Manuel, S. (1996). How word onsets drive lexical access and segmentation: evidence from acoustics, phonology and processing. In *The 4th International Conference on Spoken Language Processing*, pages 66–69.
- Hacking, I. (1983). *Representing and Intervening*. Cambridge University Press, Cambridge UK.
- Hanson, N. R. (1958). *Patterns of Discovery. An Inquiry into the Conceptual Foundations of Science*. Cambridge University Press, Cambridge UK.
- 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.
- Hurford, J., Studdert-Kennedy, M., and Knight, C., editors (1998). *Approaches to the Evolution of Language*. Cambridge University Press, Cambridge UK.
- Ishiguro, H. (1969). Use and reference of names. In Winch, P., editor, *Studies in the Philosophy of Wittgenstein*. Routledge and Kegan Paul, London.
- Kuhn, T. (2000). *The Road since Structure*. University of Chicago Press, Chicago.
- Marconi, D. (1997). *Lexical Competence*. the MIT Press, Cambridge MA.
- McDaniel, D. (2005). The potential role of production in the evolution of syntax. In Tallerman (2005).
- McGuinness, B. (1981). The so-called realism of Wittgenstein's *Tractatus*. In *Perspectives on the Philosophy of Wittgenstein*. Blackwell, Oxford UK.
- Mehler, J., Dommergues, J. Y., Frauenfelder, U., and Segui, J. (1981). The syllable's role in speech segmentation. *Journal of Verbal Learning and Verbal Behavior*, 20(3):298–305.

- Morpurgo Davies, A. (1996). *La linguistica dell'Ottocento*. il Mulino, Bologna.
- Newmeyer, F. J. (2005). *Possible and Probable Languages. A Generative Perspective on Linguistic Typology*. Oxford University Press, Oxford UK.
- Peacocke, C. (2008). *Truly Understood*. Oxford University Press, Oxford UK.
- Pustejovsky, J. (1995). *The Generative Lexicon*. the MIT Press, Cambridge MA.
- Quine, W. V. O. (1953). *From a Logical Point of View. Nine Logico-philosophical Essays*. Harvard University Press, Cambridge MA.
- Quine, W. V. O. (1960). *Word and Object*. the MIT Press, Cambridge MA.
- Quine, W. V. O. (1969). Reply to Chomsky. In Davidson and Hintikka (1969).
- Rao, P. V. S. (1993). VOICE: An integrated speech recognition synthesis system for the Hindi language. *Speech Communication*, 13(1-2):197–205.
- Read, R. and Crary, A., editors (2000). *The New Wittgenstein*. Routledge, London.
- Rendall, D., Owren, M. J., and Ryan, M. J. (2009). What do animal signals mean? *Animal Behaviour*, 78:233–240.
- Rosso, M. (1999). Introduzione. In Wittgenstein, L., *Osservazioni filosofiche*. Einaudi, Torino.
- Sperber, D. and Wilson, D. (1986). *Relevance. Communication and Cognition*. Blackwell, Oxford UK.
- Tallerman, M., editor (2005). *Language Origins. Perspective on Evolution*. Oxford University Press, Oxford UK.
- Tarski, A. (1969). *Logic, Semantics, Metamathematics. Papers from 1923 to 1938*. Clarendon Press, Oxford UK.
- Whorf, B. L. (1971). *Language, Thought, and Reality. Selected Writings of Benjamin Lee Whorf*. the MIT Press, Cambridge MA, 6th edition.
- Wittgenstein, L. (1922). *Tractatus Logico-philosophicus*. Routledge and Kegan Paul, London.
- Wittgenstein, L. (1953). *Philosophical Investigations*. Blackwell, Oxford UK.