Semantic Under-determinacy, Content and Context

Delia Belleri

No man knows what is in the mind of his friend.
(Talmudic saying)

Contents

1. Under-Specification and its Sources

Since early analytic philosophy, theorists have pointed at multiple ways in which language can be regarded as a poor means for an effective and clear expression of thought. In continuity with this tradition, contemporary authors almost unanimously recognise language as being a defective means for communication. In this thesis, I engage with the idea that the meaning of natural language sentences fails to specify the truth-conditions of their utterances—which are what speakers employ in order to describe the world or to communicate the contents of their beliefs, desires, etc. to other individuals. The following thesis is therefore taken to hold for a large number of cases:

[The Under-Specification Thesis] Given an utterance u of a sentence s in a context c, the meaning of s fails to specify the conditions in which u is true or false[1].

There are many sources of semantic under-specification. First of all, linguistic meaning can give rise to under-specification because language can be *ambiguous*, where ambiguity can be traced at more than one level of linguistic representation. There is lexical ambiguity, i.e. words can have more than one meaning: the classic example is the word "bank", which means both "financial institution" and "river side". Not only lexical items, but also syntactic forms such as phrases and sentences can be ambiguous: a good example of the former kind of ambiguity is the phrase "additional vitamin source", which means both "additional source of vitamin" and "source of additional vitamin". An example of sentential ambiguity is "Every woman loves a sailor", which means both "There is a sailor which is loved by every woman" and "For every woman, there is a sailor which that woman loves"; here the ambiguity specifically concerns the *scope*, i.e. the logical range of operation, of quantifier expressions such as "every" and "some" (also written as the indefinite article "a"). Ambiguity has always been considered as a fault of natural language insofar as ambiguous expressions cannot be assigned a meaning and a truth-value, unless the context of use provides the means for selecting one interpretation at the expense of the others.

Secondly, under-specification can arise because language can be *vague*: an expression's meaning being vague entails that there are cases in which it is unclear whether that expression does or doesn't correctly describe how the world is. An example of vague expression is the predicate "is bald". An individual with no hair on her head is clearly bald, while an individual with a hundred thousand hairs is clearly non-bald; however, there could be intermediate cases between zero hairs and one hundred thousand hairs in which it is just not clear whether we should describe the individual in question as bald or non-bald: these are typically called "borderline cases". Borderlineness is standardly considered the hallmark of vagueness. Vagueness is a fault for language insofar as it seems, at least *prima facie*, that (unless some cut-off point is arbitrarily set by the speakers) uses of sentences containing vague predicates in borderline cases cannot be evaluated for truth and falsity.

Thirdly, under-specification may come from the fact that linguistic meaning can be *context-dependent*. Context-dependent expressions typically have no stable content across contexts. The best examples of context-dependent expressions are *indexicals* such as "I", or "today". The word "I" has a stable meaning but has no stable referent, that is, the referent of "I" varies according to who the utterer is. Similarly, the word "today" has a stable meaning but it has no stable reference, for its reference varies according to the day of utterance. According to the standard semantics of

---

[1]   The failure to specify truth-conditions is best thought of as relative to focal features—such as times, places, manners, comparison classes, etc., and not as absolute.

indexicals developed by Kaplan (1989), the meaning of "I" can be described as a function that maps the context in which "I" occurs onto the speaker of that context; similarly, the meaning of "today" is a function that maps the context in which "today" is uttered onto the day of utterance. Indexicality (which is the paradigmatic case of context-dependency) makes language "faulty" insofar as, unless the values of some contextual coordinates are fixed by the interpretation (such as who the utterer is and which day is the day of utterance) it is normally impossible to assign a truth-value to utterances of the sentences that contain these expressions.

The expressions commonly recognised as context-dependent are indexicals such as "I", "you", "today", "here", "now" and demonstratives such as "this", "that", "there", "then". Authors disagree on whether the list should be extended so as to include expressions whose semantic behaviour suggests that they might be context-dependent: for example quantifiers ("every", "no", "some"), gradable adjectives ("tall", "rich", "old"), relational expressions ("enemy", "fan"), but also epistemic predicates ("know"). The dispute over whether the set of context-dependent expressions should be enlarged is highly relevant to the topic of this thesis: as it will become clear, I shall take a definite position as to the semantics of some of the expressions listed above.

The debate over whether the set containing indexicals and demonstratives should encompass more expressions reveals a whole domain of phenomena connected with semantic under-specification which however do *not* allow an immediate labelling in terms of "context-dependency", as I will illustrate later on. In what follows, I shall review those examples of this controversial domain that I deem most interesting for my purposes, which I shall call *semantically under-determined* expressions.

As it has emerged previously, all context-dependent expressions are also semantically under-specified. So, as one could expect, those expressions besides indexicals and demonstratives for which the question of context-dependency arises exhibit some form of semantic under-specification. For some of these, the question could arise whether their under-specification is a genuine one.

Some expressions are judged semantically under-specific on the basis of the fact that sentences containing them undergo contextual shifts in truth-value. Occurrences of the same expression e in different contexts c_1, c_2 … c_n give rise to different semantic values for the sentence in which e occurs. Take for example the predicate "know": in a mundane context c_1, when I consider the question of whether I have a cup of coffee in front of me, my visual experience as of a cup is sufficient to make my statement of "I know there is a cup in front of me" true. By contrast, as soon as I enter context c_2 where I engage in a skeptical exercise, it seems that my visual experience doesn't offer the required epistemic support, thus making my statement of "I know there is a cup in front of me" false. Thus, the same sentence "I know there is a cup in front of me" can be evaluated as true or false, depending on the epistemic standards that are in force in the context of use (see DeRose 1992, 1995, Cohen 1999). Similar considerations hold for the quantifier "every". Suppose that, in c_1, I assert "Everyone was accepted to the conference"; if there is no restriction over the domain of quantification of "everyone", my statement is strictly speaking false, for it is definitely not the case that every individual in the universe has been accepted for the conference; if, by contrast, in context c_2 the quantifier "everyone" is interpreted as ranging over my colleagues at the department, it turns out as true iff all of them have been accepted for the conference (on quantifier domain restriction, see Neale 1990, Stanley & Williamson 1995, Lewis 1996, Gauker 1997, Stanley & Szabó 2000). The fact that sentences containing "know" or "every" can give rise to shifts in truth-value across contexts is indicative of under-specification. However, I urge that the evidence is not conclusive. The reason is that the possibility remains open that statements like "I know there is a cup in front of me" and "Everyone was accepted to the conference" express constant, determinate contents across contexts, namely *<that I know* (in the only sense of "know") *that there is a cup in front of me>* and *<that everyone* (in the unrestricted sense) *was accepted to the conference>*. What varies across contexts, one could contend, is just the truth-value of these utterances (see Stanley 2005a, Hawthorne 2004 concerning "know") or at least the truth-value of a further content which is

"suggested" or "implicated" by the fact of making those statements in either context c_1 or c_2 (see Cappelen and Lepore 2005). What I am saying is therefore that, for expressions like "know" or "every", contextual shifts in truth-value do not *ipso facto* constitute evidence of their semantic under-specification, because the possibility is not ruled out that "know" and "every" respond to an *Invariantist semantics*, which assigns a stable and determinate content to statements that contain those expressions and explains the utterances' variations in truth-value by appealing to either relative truth or to pragmatically suggested contents[2].

Similar considerations apply to sentences like "Jack and Jill got married", "I've had breakfast", "It will take time to get there", "You are not going to die". In all these cases, it has been argued that the meaning of these sentences under-specifies, in some sense, the truth-conditions of its utterances (see Carston 1988, Recanati 1989). For example, the sentence "Jack and Jill got married" is true just in case Jack got married and Jill got married; however, speakers usually do not mean just that; what is usually meant (and understood) is a content that is true just in case Jack and Jill got married *together*. Similarly, by uttering "I've had breakfast", one usually means something which is true only if one has had breakfast within a reasonably short span of time—that very morning, for example—and not just in case one has had breakfast at least once in one's life. By claiming "It will take time to get there", the normally intended content is true just in case it will take *more time than expected* to get in the salient place, for the literal meaning of the sentence is a sheer triviality, not worth stating in normal circumstances. Finally, one may not want use "You are not going to die" while meaning that, literally, the addressee is not going to die at all—for this would be a patent falsehood. As one can see, the meaning of each of these sentences expresses a content which is not the content that speakers may usually mean while uttering these sentences—whence the contention that these sentences are "under-specific". On closer inspection, though, the under-specificity of these sentences may not be a semantic one at all, for each of them expresses *an entirely evaluable content*. With respect to the semantics of each of these sentences, it makes perfect sense to adopt an Invariantist stance, having it that they express a determinate and constant content across contexts. The only fault of these sentences is that of expressing semantic contents which—albeit determinate and invariant—are not identical to what speakers usually mean by uttering them. This however, doesn't look like a genuine, "deep" semantic under-specification, but as a "shallow", *pragmatic* under-specification.

Positing the under-specification of words like "know", "every", "and", but also of sentences like "I've had breakfast", "It will take time to get there", "You are not going to die" is a way of explaining some sets of intuitions, especially those concerning cross-contextual "shifts in truth-value" and the gap between what the sentence says and what speakers usually mean; however, since it at least *makes sense* to go Invariantists about the semantics of these expressions, the case for their "deep", *semantic* under-specification is not entirely convincing. After all, theirs could be just a "shallow", *pragmatic* under-specification, that is, these expressions may truly express a determinate content across all contexts, but this content may not be what speakers usually *mean* to communicate.

2. Semantic Under-determinacy
Besides these potential examples of "shallow" under-specification it is possible, however, to

---

[2] I take Grice (1967/1989) to have made an Invariantist point about e.g. the semantics of logical constants in natural language such as "but", or "therefore": in his view, the truth conditions of sentences like "He is rich but generous" and "He is an Englishman, he is therefore brave" are the same across both informal conversations and formal discourse, namely they are the truth-conditions of a conjunctive sentence; what changes is what, in informal discourse, the use itself of "but" and "therefore" suggests over and beyond the meaning of the sentence itself. Kripke (1977) defended the same position concerning the truth-conditions of sentences containing incomplete definite descriptions, by distinguishing between "semantic reference" and "speaker's reference". Concerning the semantics of knowledge predicates, even more radical versions of Invariantism than those mentioned in the text are defended by Brown (2005) Black (2008) and Turri (2010).

identify a group of expressions whose "deep", semantic under-specification can be shown more compellingly. These expressions are such that (i) they give rise to contextual shifts in truth-value; (ii) going Invariantist about them really makes little, if any, sense; (iii) there is a straightforward intuition about their semantic defectiveness, independently of contextual shifts. These expressions, that I shall characterise as *semantically under-determined*, are going to be the focus of the present research. Note in passing that points (i)-(iii) are *not* meant to state either necessary or sufficient conditions for what I call semantic under-determinacy, but merely a "singling-out procedure", i.e. a practically useful set of instructions for recognising typical cases of semantic under-determinacy as opposed to other cases of under-specificity. In what follows, I will confine myself to a presentation of some examples of semantically under-determined expressions, with only a sketchy discussion of their potentially problematic aspects. My primary interest at the moment is *to make sense* of the idea of semantic under-determinacy itself, independently of the solutions that might be elaborated in order to frame it semantically. In introducing examples of semantic under-determinacy, I will be partially following a recent introductory survey provided by Bach (2012). Let us consider a sentence like

(1) It's raining.

This sentence certainly undergoes variations in truth-value with relevant variations in context. For example, if I utter (1) while talking about Paris, in which rain is falling, my utterance of (1) shall be true. If I utter the same sentence while talking about Venice, where the sun is shining, the same sentence shall be false. The sentence "It's raining" therefore meets requirement (i) by being subject to contextual shifts in truth-value. This is already indicative about the semantic under-specification of (1). However, one may argue, it is always possible that the predicate "to rain" has an Invariantist semantics: it determines a stable, full-blown content across contexts, and what is subject to variations is the content of the speech act performed in each different context (see Cappelen and Lepore 2005). This kind of Invariantism about "It's raining" I take to be seriously problematic. First of all, it is not sensible to suppose that an utterance of the sentence "It's raining" expresses a content which is true iff it's raining *tout court.* There is no such thing as raining *tout court*—rainings are invariably located at some place. Secondly, even if there was a determinate content for (1), perhaps *<that it's raining somewhere>*, it's not clear how the sentence "It's raining" could express exactly this content, for it seems that that content is not strictly speaking semantically encoded: if one only considers the linguistic meaning of the sentence's components and the way they are combined, it seems that the result is simply the content *<that it's raining>*. For these reasons, going Invariantists about (1) would make little sense, as stated at point (ii). Finally, and as stated by point (iii), we do not need contextual shifts to make the under-specification of (1) apparent, for in the appropriate conditions we can have a direct intuition of its defectiveness. This is confirmed by the fact that, if one were to utter (1) out of the blue and with no apparent reason, it would be natural for a hearer to ask for specification with questions like "Where?", "In which place?" and so on. These questions I take to be stimulated by the sentence's inherent lack of specificity about the utterance's truth-conditions.

These considerations strongly suggest that (1) is *semantically under-determined*. Do they make for the thought that "to rain" is context-dependent, too, in the sense that it involves some covert form of indexicality? Here the point is debatable; although theorists have argued for the context-dependency of "It's raining" (see Stanley 2000), the arguments that have been worked out so far have not encountered a unanimous consensus (see e.g. Neale 2007). Moreover, the contextual variation in truth-value of "It's raining" may not depend on a defectiveness of content, but it could depend on a relativity of the truth-value of (1) to a parameter corresponding to a suitably cut-up "situation" or world-region (see Barwise and Perry 1983). In other words, no absolutely compelling case has been made yet to settle the context-dependency of (1). For the moment, I will be content with taking this

"sociological" fact as a reason for claiming that a semantically under-determined sentence like "It's raining" may be so in a way that has nothing to do with context-dependency (in the form of covert indexicality). More cases of semantic under-determinacy that prove similar to (1) are:

(2) It is windy.
(3) It is noisy.
(4) It is eerie.

The list can be extended with expressions which are semantically under-determined insofar as, as Bach has it, they can be used without what in the linguistics jargon are called *complements* (see Bach 2012). Below are two examples:

(5) Jill is ready.
(6) Lynn is late.

Here it's not difficult to ascertain that, for example, sentence (5) responds to requirements (i)-(iii). To begin with, point (i) is met, for (5) changes truth-value across relevantly different contexts: if I utter (5) while talking about Jill's preparing her Spanish test, and Jill is absolutely unprepared for the exam, then my utterance is false; if I utter (5) while talking about Jill's getting ready to go to the club and she is absolutely ready for that, then my utterance is true. Point (ii) is also met, for Invariantism about "being ready" barely makes sense: first of all, an utterance of (5) cannot be true iff Jill is ready *tout court*, for there is nothing like being ready *tout court*: the predicate "being ready" at least implicitly suggests that readiness is instantiated with respect to a purpose. Secondly, even if (5) had an invariant content like *<that Jill is ready for something>*, it's not clear how the sentence "Jill is ready" could encode that content, for it seems that such a sentence does not strictly speaking have that meaning: indeed, by the lights of formal semantics, the linguistic meaning of (5) is simply *<that Jill is ready>*. Finally, ordinary speakers have the chance to directly intuit the semantic defectiveness of (5): if one were to utter "Jill is ready" out of the blue and with no evident motivation, the hearer would be entitled to prompt the speaker with questions like "For what?", "For which purpose?". This I take to reflect the fact that the truth-conditions of (5), such as they are determined by semantic conventions, do not tell how the world should be for an utterance of (5) to be true. These considerations again make for the idea that "Jill is ready" is semantically under-determined.

At the same time, it's not clear that we should draw the additional conclusion that "Jill is ready" also is context dependent in the way of indexical expressions. There is no clear and compelling evidence that "being ready" is context-dependent in the same way as "I" or "this". Moreover, data concerning the semantic variability of (5) could receive an alternative explanation, which doesn't involve hidden indexicality but rather involves the postulation of a truth-predicate which is relative to a suitable parameter in the circumstances of evaluation (see MacFarlane 2007a, Gauker 2010) or to a situation (see Corazza and Dokic 2010). Once again, at least if we content ourselves with the "sociological" fact that the context-dependence of "ready" is still a matter of controversy between scholars, it seems that (5) exhibits semantic under-determinacy but not, for that matter, context-dependency (such as a form of covert indexicality).

The same reasoning as above could be applied to the semantic defectiveness that characterises gradable adjectives, such as 'fat', 'tall', 'old', 'fast', 'rich':

(7) The cat is fat;
(8) Lisa is tall;
(9) Fido is old;
(10) That car is fast;

(11) Naomi is rich;

Consider for example sentence (8). Once again, it seems that points (i)-(iii) are all met. As provided by (i), the sentence "Lisa is tall" varies its truth-value with relevant variations in context. If the salient comparison class is 8-year-old children and Lisa is an extraordinarily tall 8-year-old girl, then my utterance is true; if the salient comparison class is adult basketball players, then my utterance is false. Also, accordingly with (ii), going Invariantists makes little or no sense. An utterance of (8) cannot express a content that is true just in case Jill is tall *tout court*, for there's nothing like being tall *tout court*: the notion of "being tall" at least implicitly suggests the presence of a comparison class, or a standard. At the same time, granting (8) expressed a content like *<that Jill is tall with respect to some comparison class>*, it remains to be explained how (8) could semantically encode exactly that content since, by the lights of formal semantics alone, it seems that the linguistic meaning of (8) is simply *<that Jill is tall>*. Finally, as in (iii), judgements as to the semantic defectiveness of (8) come spontaneously in the appropriate circumstances, for were (8) to be uttered just out of the blue, the hearer could legitimately reply with a request for completion like "For what?", "With respect to whom?". This would plausibly result from (8) being directly understood as semantically under-specific.

Thus again, we seem to have semantic under-determinacy. But do we have elements for concluding that there is context-dependency, too, perhaps in the form of hidden indexicality? Once again, the arguments that have been designed so far in favour of the context-dependency of "tall" and gradable adjectives in general (Stanley 2000) have not been universally accepted. An alternative explanation is always available, which consists in relativising the truth of (8) to some parameter of evaluation (see MacFarlane 2007a, 2009, Richard 2004, 2008) or a situation (see Corazza & Dokic 2007, 2010). Thus, we may conclude provisionally, and once again on the basis of "sociological" considerations about the state of the art, that the semantic under-determinacy of (8) may have nothing to do with forms of context-dependency such as hidden indexicality, unless stronger and more persuasive arguments are provided.

The list could be extended with so-called *relational* predicates such as "local", "enemy", "fan", "neighbour", "immigrant". Here the semantic under-determinacy of the sentences becomes particularly apparent for, as Bach notes, sentences like the following barely make sense unless a particular perspective is assumed as salient:

(12) Oliver is a neighbour.
(13) Oscar is a fan.
(14) Osama was an enemy.

One can easily realise that (12)-(14) positively respond to the "singling-out procedure" suggested above: (i) they are subject to contextual shifts in truth-value; (ii) it makes little sense to go Invariantists about their content and truth-value; (iii) direct intuitions of their under-specification can be obtained whenever an utterance of (12)-(14) doesn't seem to presuppose any particular *relatum* for the predicates "being a neighbour of", "being a fan of", "being an enemy of". Theorists have advanced the idea that relational predicates such as "enemy", "fan" are covertly indexical, i.e. they contain implicit argument places in their deep syntax, such that they can be either assigned a value or be bound by an operator (see Stanley 2000; Partee 1989 considers and discusses this option, even though she eventually discourages it). Theorists tend to disagree less about cases like (12)-(14), because the implicit structure of predicates such as "being an enemy" leaves little room for solutions that dispense altogether of relational configurations[3]. Still, as I will clarify in the

---

[3] For example, speaking about relational predicates such as "local" and "immigrant" Corazza acknowledges that they are most intuitively characterised as relations, and thus as context-dependent terms: "The intuitive way to understand [utterances that contain "local" or "immigrant"] is that the predicate itself is context sensitive and, thus,

following sections, I believe that syntactic commitments are not strictly speaking necessary here. One can view the relational configuration of predicates such as "enemy", "fan" etc. as being implemented at the *conceptual/semantic* rather than syntactic level, where syntactic structure and conceptual/semantic structure differ mainly in that the first is part of our linguistic competence, while the second certainly interacts with linguistic competence, but it is also part of a more general system of competences that include our knowledge of the world, and in general of our social, cultural environment.

A further group of expressions that could be labelled as semantically under-determined is the group of what we may call *perspectival* predicates, i.e. predicates whose application requires that a point of view or perspective is either explicitly expressed or implicit but salient in the conversational setting. These include predicates that express spatial relations, such as "being left", "being right", "being nearby" "being outside", "being above", etc. but also predicates that express response-dependent properties, such as "being edible", "being funny", "being terrifying", "being exciting". These predicates will also include so-called *predicates of personal taste*, which express properties that items possess with respect to a subject's taste standards, such as "tasty", "delicious", "disgusting". Sentences containing these predicates will count as semantically under-determined to the extent that, as long as one remains neutral as to their semantics, intuitions tell us that (in many if not all of their usages) they fail to specify the appropriate relativisation for the predicate at issue at the level of utterance truth-conditions—in other words, they fail to specify *for whom* a certain item is tasty, delicious, disgusting, etc[4].

(15) The post office is on the left;
(16) Naomi went to a beach nearby;
(17) The taxi is waiting outside;
(18) Some roots are edible.
(19) That joke is funny.
(20) That movie is terrifying.
(22) Ice cream is delicious;
(23) Rotten meat is disgusting.

Once again, the singling-out procedure previously articulated gives positive results. In each of these cases: (i) utterances of the same sentence in a relevantly different context give rise to a truth-value variation; (ii) it makes little or no sense to go Invariantist with respect to the content expressed by each sentence: for example, there seems to be no absolute way for anything to be on the left, nearby, or outside; nor there seems to be anything like being edible *tout-court* (which is certainly not to be equated with "being edible by every sentient being"), or being funny, terrifying, delicious *tout court.* All these expressions seem to at least suggest an underlying relational semantic structure. Even conceding (15)-(23) expressed a determinate content, as e.g. *<that the post office is on the left of something>*, *<that ice cream is delicious for someone>* it is not clear how these contents could be semantically encoded in the sentences at interest: a formal, compositional approach to the

_____

that the context sensitivity of the utterance is triggered by the fact that the predicate is relational. The natural way to understand these predicates is to view them as two-places predicates: a location is local *vis-à-vis* someone, *local(p, x)*; a person is an immigrant *vis-à-vis* some country, *immigrant(c, x)*. The dictionary registers that 'local' means "existing in or belonging to the area where you live, or to the area you are talking about," that 'immigrant' means "a person who settles as a permanent resident in a different country," that "a foreigner is someone who belongs to a country which is not your own," and so forth." (Corazza 2007: 433).

4   I should warn the reader that that the fact that predicates like "tasty", "disgusting" are putatively semantically under-determined should *not* be conflated with the question of their truth-conditions being intrinsically relative. That the two problems do not coincide can be seen in theories like Contextualism, which holds that the truth-value of utterances of "This ice cream is tasty" is absolute, while at the same time admitting that the sentence "This ice cream is tasty" doesn't specify the truth-conditions of its utterances, in that it doesn't specify *for whom* the ice cream is tasty.

semantics of these sentences really delivers simple and "minimal" linguistic meanings such as *<that the post office is on the left>*, *<that ice cream is delicious>*. Finally, as provided by (iii), in appropriate conditions, an utterance of each of (15)-(23) could prompt requests for completion of the form "Left of what?"; "Edible for whom?", "Delicious for whom?"; all these requests of supplementation most plausibly are triggered by the sentence's inherent semantic defectiveness. Thus, we have a basis for affirming that examples like (15)-(23) are semantically under-determined. The issue of the context-dependency of spatial, response-dependent and personal-taste expressions is once again a matter of controversy. Some argue that predicates expressing spatial relations are context dependent (see Partee 1989), while some others argue that there is some context-dependence, which nevertheless doesn't reduce to sheer indexicality (see Hawthorne 2006). Some have argued that response-dependent and taste predicates are really context-dependent (Cappelen and Hawthorne 2009, Lopez de sa 2008, Sundell 2011), while others have preferred an explanation in terms of relative truth (Kölbel 2002, Lasersohn 2005, MacFarlane, MS.). The debate is still open. For all we know, and given the ongoing debate among semanticists, the semantic under-determinacy of (15)-(23) could be totally disconnected from context-dependence (in the form of covert indexicality).

The last group of predicates which, if I am right, should be subsumed under the semantic under-determinacy category, is that of *colour predicates*. It's not immediately clear that colour predicates should be semantically defective in any way. However, reflection concerning the conditions of application of any expression denoting a chromatic property reveals the contrary. I take colour predicates to range from predicates that attribute a certain shade of colour to predicates that attribute also other properties that pertain to an object's visual appearance, such as being clear, being dark etc.

(24) The leaves are green;
(25) The surface is red;
(26) The room is dark;

Sentences containing colour predicates certainly undergo contextual shifts in truth-value, as in (i): Travis notoriously argues for the variability in truth-value of (24) across different contexts. Pia, who has painted green the leaves of her red Japanese maple utters "The leaves are green" while talking to a photographer interested in green-looking items: on this occasion, she says something true. Later on, she utters (24) while talking to a botanist interested in naturally green items: on this occasion, she says something false. That going Invariantists about "green", or "red" or "dark" is really a bad idea may not be immediately evident. However, a little reflection helps to recognise that there is nothing like being green, or red, or dark *tout-court*. A colour property is generally attributed to an object in some respect: greenness is invariably greenness in some part, under certain light conditions, from a certain distance. Also, attributions of colour properties depend much on contrast items: a pigeon's plumage may be described as clear if contrasted with a raven's, but it may be described as dark if contrasted with a swan's. All these considerations help make sense of the idea that colour attributions cannot be absolute—they invariably involve (at least implicitly) a form of relativity to some respect. One could argue that a predicate such as "green" means "green in some respect", thus making it that (24) expresses the content *<that the leaves are green in some respect>*. However, once again I don't see how a proponent of this view could positively argue that exactly that content is strictly speaking encoded by (24): as far as the compositional semantics of the sentence goes, it really seems that the content linguistically expressed by (24) is simply *<that the leaves are green>*. Going Invariantists about colour predicates is thus not a good idea, accordingly with point (ii). The semantic defectiveness of sentences like (24)-(26) can be intuitively judged in the appropriate conditions, as provided by (iii): when a hearer doesn't have sufficient background information, she is entitled to ask for completion with questions such as "Green in what respect?",

"Dark with respect to what?".

So sentences containing colour predicates do suffer from semantic under-determinacy. Are they context-dependent, too, perhaps because they involve some hidden indexicality? Arguments purported to prove this (see Szabó 2001, Rothschild & Segal 2009) have been addressed as exceedingly complicating the semantics of these expressions (see Recanati 2010, Hall 2008). As I will argue throughout this thesis, that the attribution of colour predicates suggests some hidden complex structure may not be a question of syntactically articulated positions, but rather a question of implicitly assumed *conceptual/semantic* structures. Once again, with colour predicates we have an example of expressions that exhibit some under-specificity, but that do not immediately suggest a covert indexicality, as it emerges from the fact that scholars haven't reached an agreement as to their syntactic structure.

To sum up, in this section I have aimed at giving the reader an idea of what semantic under-determinacy is. I have provided a (provisional) "singling-out procedure" for semantically under-determined expressions, which contemplates three points: (i) semantically under-determined expressions give rise to cross-contextual shifts in truth-value; (ii) it makes little if no sense to go Invariantists as to the semantics of these expressions; (iii) intuitions as to their semantic defectiveness can be triggered independently of any contextual shift. Expressions which meet points (i)-(iii) have a good chance of being semantically under-determined; an additional, indirect piece of evidence that this might be the case is given by the fact that scholars haven't reached an agreement yet concerning the underlying syntactic structure of these expressions, in that they haven't clearly established whether these expressions are, for instance, context-dependent, perhaps because they involve hidden forms of indexicality.

3. The Under-articulation Account

The preceding sections have been devoted to tracing the boundaries of the phenomenon on which I wish to concentrate in my research: semantic under-determinacy. So far I have reached the following, rather provisional result: semantic under-determinacy is a kind of semantic under-specification, which manifests itself along the lines traced by (i)-(iii), and which doesn't immediately strike as an example of context-dependence, as for instance a case of hidden indexicality. I have been concerned with offering a good number of examples of semantic under-determinacy, the primary aim being that of helping the hearer to make sense of the notion itself. For ease of exposition, I repeat here some of the examples with which I have illustrated the notion of semantic under-determinacy:

(1) It's raining;
(5) Jill is ready;
(8) Lisa is tall;
(17) The taxi is waiting outside;
(24) The leaves are green.

So far, I have merely hinted at the fact that semantic under-determinacy is different from ambiguity, vagueness, indexicality, but I haven't entered into a detailed argument about how this difference should be traced and motivated. In **Chapter 1** of the thesis, my initial concern is precisely that of differentiating semantic under-determinacy from these phenomena. Indeed I maintain that, if one wishes to gain full comprehension of the phenomenon of semantic under-determinacy, one should be able to sharply demarcate it as opposed to other phenomena of semantic under-specification. The problem I shall deal with in **Chapter 1**, which I will call "The Demarcation Problem", may be thus formulated:

[The Demarcation Problem]: Is there a feature which distinguishes semantic under-determinacy

from other phenomena of semantic under-specification?

Answering the *Demarcation Problem* requires pointing at at least one feature which semantically under-determined expressions have, but which other expressions do not have, which may enable one to trace a significant distinction between the former and the latter. Note that dealing with this problem doesn't require that one identifies necessary *and* sufficient conditions for semantic under-determinacy—the problem may be successfully dispensed with even by providing merely necessary conditions. In **Chapter 1**, I argue that the feature which distinguishes semantic under-determinacy from the rest of the other semantically under-specific expressions is that it is *explainable solely in terms of under-articulation*, that is to say in terms of lack of linguistic material only. In the present section, I will anticipate the main arguments I will provide in favour of this thesis.

In **Chapter 1** I draw attention on the fact that some theorists have recommended to resist the temptation to "domesticate", i.e. to reduce, semantic under-determinacy to any other familiar phenomenon, such as the already mentioned ambiguity, vagueness, indexicality. This I refer to as a recommendation to follow a "no-domestication" policy.

Let us start with the temptation to claim that semantic under-determinacy is reducible to *ambiguity*. I take it that the necessary conditions for ambiguity at the lexical level are that (1) there be a single lexical entry to which more than one meaning is assigned by the dictionary; (2) the number of the meanings, although it can increase or decrease as time goes by, has to be definite, for reasons connected to language acquisition and memory. If we move to the syntactic level, the condition becomes that there be a single combination of lexemes which can be associated with a definite number of phrasal/sentential tree-structures. Let me now consider one of the sentences which we have characterised as semantically under-determined: "Jill is ready". *Prima facie*, this sentence seems to be associated with *one single* tree-structure, thus violating the necessary requirement for syntactic ambiguity. Moving to the lexical level, one could surmise that the lexical entry "ready" is associated with more than one meaning *ready$_1$, ready$_2$, ready$_3$ … ready$_n$*, which gets selected according to the context of utterance. To this idea one can object that, since one can be ready for indefinitely many purposes, the word "ready" will have to be ambiguous between indefinitely many meanings. This goes against the necessary requirement (2) for ambiguity. In sum, there are *prima facie* good reasons for resisting a reduction of semantic under-determinacy to ambiguity (see also Travis 1997, Bezuidenhout 2002).

One could sympathise with the view that semantic under-determinacy is reducible to *vagueness*. Even though the matter of what vagueness is is itself under discussion let us assume, in line with the mainstream conception, that a necessary feature of vagueness is the obtaining of borderline, or unclear cases (e.g. of baldness, tallness, fatness, etc.). Now one should definitely concede that expressions like "It's raining", "Jill is ready", "Lisa is tall", "The leaves are green" do contain predicates that give rise to borderline cases—after all, there are borderline cases of raining, as well as borderline cases of readiness, tallness and greenness (the latter two are indeed paradigmatic cases of vagueness!). However the particular defectiveness that makes these expressions semantically under-determined is *not* connected to their being vague. To illustrate, consider an utterance of "It's raining" which one assesses as not determinately true or false (maybe because drops of water are falling from the sky at too large a distance from each other); now I figure that, in order for one to be able to say that the utterance doesn't have a definite evaluation in the first place, one needs to figure out the location where the phenomenon is taking place—Paris, Venice, London … . That is, in order to judge about the vagueness of a statement of "It's raining" one must previously resolve its semantic under-determinacy as to the location of the rain. Similarly, one could envisage an utterance of "X is green" which is clearly true, because x is a clear case of greenness. However, this doesn't erase the fact that the sentence uttered suffers from semantic under-determinacy, in that it doesn't tell in what respect x is green. This seems to make a case for the idea that semantic under-determinacy is not reducible to vagueness (see also Travis 1997, Bezuidenhout 2002).

What about the idea that semantic under-determinacy is reducible to indexicality? In the preceding sections, I mentioned the fact that there's quite a lively debate with regard to this issue. Some theorists starkly defend the idea that certain semantically under-determined expressions, as e.g. "It's raining" (Stanley 2000) or "The leaves are green" (Szabó 2001) are really covertly indexical. I have briefly mentioned that other theorists oppose this attempt to reduce semantic under-determinacy to indexicality by appealing to various ways of relativising the truth-predicate. The fact that no account has clearly prevailed in the debate already suggests that a reduction of semantic under-determinacy to indexicality is not straightforward. This, however, is nothing more than acknowledging a *sociological* fact: what is needed is an argument that shows the point at the conceptual/theoretical level, not merely at the sociological one.

Just like indexicals, semantically under-determined expressions fail to specify the content-in-context of their utterances: for instance, the linguistic meaning of "I wear glasses" fails to specify the content of its utterances, for the simple reason that "I" needs to be assigned a referent in context. In the same vein, one could say that the linguistic meaning of "It's raining" or "Jill is ready" or "The leaves are green" fails to specify the content-in-context of their utterances. The content of these sentences needs to be supplemented in some way: for example, one needs to specify the location of the rain, the purpose of Jill's being ready, the respect in which the leaves are green, etc. Is this equivalent to the assignment of a reference to "I", or even to "here", "there", etc.? My claim is that there is a difference between how indexicality is resolved and how semantic under-determinacy is resolved. Let me preliminarily distinguish between *pure* indexicals and *spurious* indexicals. The former are indexicals whose reference is fixed independently of the beliefs of the speaker: in this sense, "I" is certainly a pure indexical. The latter have their reference fixed also as a function of the speaker's beliefs and intentions: for instance, the reference of "here" can vary in its width depending on what portion of space the speaker intends to refer to with her use of "here". Given this distinction, I shall say that at least *pure* indexicals are mechanisms of contextual reference-assignment that are designed for being *determinately saturated*. That is, at least pure indexicals are designed in such a way that, once the term's occurrence is assigned a referent according to the value taken by a certain parameter in a certain context k, the under-specification of the expression is *ipso facto* resolved, in the sense that its content becomes as specific as it may be, even in any further circumstance c in which the utterance may be re-considered. For example, "I" is construed in such a way that, once the reference of "I" in "I wear glasses" has been assigned in a context k, the resulting content being *<that X wears glasses>*, the content-in-context of the pronoun "I" is as specific as it may be, no matter the circumstances c_j in which the original utterance may get reconsidered. By contrast, if a semantically under-determined expression gets supplemented with extra content in a certain circumstance c_k, it is not obvious that the under-specification is resolved for all the circumstances $c_j \ldots c_i$ in which the utterance could be re-considered. For instance, suppose that the content of "It's raining" is supplemented in the circumstance c_1 so as to result in *<it's raining in Paris>*. In some other circumstance c_2, the following question could arise: is it raining in Paris if it is raining only in the center of the city? In circumstance c_3, one could wonder: is it raining in Paris if it is raining over many scattered spots of the town? I envision that the content *<that it's raining in Paris>* obtained in c_1 may turn out as not responding to the demands that may arise in circumstances c_2 and c_3, where more precision is required. Let us call this phenomenon of potential insufficiency of a contextually enriched content to respond to the precision demands that arise in other circumstances "Reiteration Effect". Thus, semantic under-determinacy is set apart from (at least pure) indexicality by the fact that (pure) indexicality resolution is set up in such a way that it's expected not to generate further under-specificity (in other circumstances), while semantic under-determinacy resolution may generate further under-specificity (in other circumstances), thus giving rise to a "Reiteration Effect".

Having argued for a distinction between semantic under-determinacy and, respectively, ambiguity, vagueness and indexicality, I believe that a crucial step should be taken. Besides distinguishing

semantic under-determinacy only in negative terms—like in the case of ambiguity and vagueness—or in terms of the *effects* it produces—such as the Reiteration Effect—the theorist should also try to trace the distinction on the basis of what *causes/explains* semantic under-determinacy. In this way, I believe that the answer to the *Demarcation Problem* would be deeper and theoretically more satisfactory.

The answer I suggest for the *Demarcation Problem* starts from the following consideration: it can be argued that semantic under-determinacy obtains because (and only because) the sentences that we use are *under-articulated*. I shall embrace the following claim:

[The Under-Articulation Claim] The meaning of a sentence *s* under-determines the content and truth conditions of its utterance *u* because (and only because) *s* is under-articulated.

A sentence *s* is under-articulated iff, when an utterance *u* of *s* is performed, another sentence *s'* could have been uttered instead of *s*, which articulates more linguistic material than *s* and has the same truth-conditions as *u* (which I take to be fixed by the utterer's intentions). In other words, a sentence *s* is under-articulated iff the truth-conditions expressed by the speaker in a context *c* with an utterance *u* could have been expressed by a more articulated sentence *s'*.

It is arguable that under-determined sentences turn out as being so because (and only because) they are under-articulated, i.e. they are "too brief" to exhaustively express the intended content-in-context of their utterances. So for example if, talking about Paris, Eva utters "It's raining", we may say that the sentence that she utters is under-determined because (and only because) Eva, in the same context, could have used the sentence "It's raining in Paris", which articulates more linguistic material and expresses the same truth-conditions of her actual utterance of "It's raining" as performed while talking about the weather in Paris. Analogously, if Eva utters "Jill is ready" meaning *<that Jill is ready to take the exam>*, we may say that the sentence that she uses is under-determined because (and only because) she could have used a more articulated sentence in order to express the same content, for example "Jill is ready for the exam". Similarly, if Eva utters "The leaves are green" meaning *<that the leaves are green on the outer surface>*, we may say that the sentence she used is under-determined because (and only because) she could have used a more articulated sentence like "The leaves are green on the outer surface" in order to express the same content.

The answer to the Demarcation Problem I would like to advance is therefore the following: what distinguishes semantic under-determinacy from other phenomena like ambiguity, vagueness and indexicality is that it is *explainable solely in terms of under-articulation.* Let me elaborate on this proposal. It seems fair to claim that ambiguous and vague expressions are such *because* of under-articulation: after all, were the speaker to use more words, she would have the chance to avoid both ambiguity and vagueness. Moreover, indexicality as well may be caused by a form of under-articulation: after all, one could always use her full name instead of the pronoun "I", thus avoiding the under-specification proper of indexicality. Even though all this seems correct at first glance, my view is that the claim is not sustainable that these expressions suffer from ambiguity, vagueness or indexicality *only because* of under-articulation. Some additional mechanism is required in order to explain the specific semantic under-specification they suffer from: for ambiguity, there is the fact that more than one meaning is associated with the same word; for vagueness, there is the fact that there are cases of unclear application of predicates; for indexicality, there is the fact that some words can be used as context-sensitive devices of reference. The only kind of semantic under-specificity for which it seems possible to argue that it is generated *exclusively* by lack of articulation is semantic under-determinacy (presumably, both in its "deep" and "shallow" varieties, though the focus here is just on the former). To recapitulate, the answer I suggest for the *Demarcation Problem* is the following: the feature that distinguishes semantic under-determinacy from other phenomena of semantic under-specification such as ambiguity, vagueness and indexicality is that, unlike these

phenomena, semantic under-determinacy is *explainable in terms of under-articulation only*.

I believe that embracing an under-articulation account as a way of answering the *Demarcation Problem* marks an advancement in our understanding of the peculiar kind of semantic under-specification of sentences like "It's raining", "Jill is ready", "The leaves are green", etc. These sentences have been troubling for philosophers of language on the account that their semantic under-specification couldn't be clearly assimilated to any familiar type of under-specification, such as ambiguity, vagueness, indexicality etc. My suggestion in the present work is that semantic under-determinacy should be thought of as a *sui generis* phenomenon, whose boundaries can be effectively captured by emphasising the idea that it is a phenomenon explainable in terms of under-articulation only.

4. Communication despite Under-determinacy

After having faced the issue of the demarcation of semantic under-determinacy as opposed to other phenomena such as ambiguity, vagueness, indexicality, it is time to move to another pressing problem which relates to semantic under-determinacy. The problem is how, given this kind of under-determinacy, successful communication is possible. How is it that speakers who use under-determined sentences manage to be understood by their hearers? The problem is stated by considering the point of view of the potential listener of one of the sentences that have been surveyed so far. What is to be explained is the general mechanism by which the listener of an utterance of "It's raining", "Jill is ready", "The leaves are green" etc. manages to understand what the utterer means. Since the issue specifically concerns the *comprehension* of under-determined utterances, it could be formulated as follows:

[The Comprehension Problem] What is the general mechanism by which the listener of an under-determined sentence *s* acquires comprehension of what the speaker means by uttering *s*?

A variety of answers is compatible with the Comprehension Problem. Assessing whether there is one answer that fares better than the others is what I focus on in **Chapters 2**, **3** and **4** of this thesis. I will isolate the answers that could be given by some of the most popular philosophical approaches to utterance comprehension; I will discuss and criticise each of these answers, and finally I will propose a new and hopefully more perspicuous solution, which makes use of the notion of "conceptual constraints". In this section, I shall provide a summary of the contents of **Chapters 2**, **3** and **4**, in which I consider different ways of tackling the Comprehension Problem.

In **Chapter 2**, I consider a Radical Contextualist approach to utterance comprehension. The central claim of Radical Contextualism is that the truth-conditional aspects of an utterance, equated with what a speaker "says" (as opposed to what she "implies") are determined to a significant degree by *free pragmatic processes*, i.e. by pragmatic processes that do not respond to any syntactic constraint. Such processes include free enrichment, strengthening and loosening of meaning (see Carston 2002, Recanati 2004), semantic transfer (see Nunberg 1995) and so on. The Radical Contextualist answer to the Comprehension Problem will thus be phrased as follows:

[The Radical Contextualist Response] Free pragmatic processes are necessary in order to achieve comprehension of utterances of semantically under-determined sentences.

The Radical Contextualist Response entails that, if no pragmatic process can be carried out by the hearer, perhaps because she has no access to the contextual information that would allow her to pragmatically supplement the sentence's defective meaning, then the hearer will not understand (i.e. know the truth-conditions of) an utterance of a semantically under-determined sentence[5]. I argue

---

[5]   Note that the same does not hold for what I call "shallow" under-specificity: in the shallow case, comprehension is still possible, by the lights of the Radical Contextualist, even in the absence of pragmatic effects. For instance, if a

that this is not true in general. Hearers can and do understand utterances of under-determined sentences, even if they do not have access to contextual information. The whole issue obviously turns on what one means when one talks about "understanding" or "comprehension".

My criticism of the Radical Contextualist Response takes as its starting point the remark that the Radical Contextualist Response implicitly assumes a notion of comprehension which I call RC-comprehension. As proponents of the Radical Contextualist approach seem to assume, RC-comprehension is knowledge of the truth-conditions of an utterance in the specific context in which the utterance is made (see Sperber & Wilson 1986/1995, Carston 1999, Recanati 1989). Thus, a hearer acquires RC-comprehension of an utterance of "It's raining" just in case she has the information available that the speaker is talking about e.g. Paris, and she uses this information in order to derive that the utterance is true just in case it's raining *in Paris*. If we interpret "comprehension" as RC-comprehension, then it correctly follows that a precluded access to the context in which the utterance is performed implies a lack of comprehension. This much can be granted.

However, RC-comprehension is certainly not the only notion of comprehension available. I introduce the notion of *Intuitive Comprehension,* which is knowledge of the truth-conditions of any arbitrary utterance of a sentence in any arbitrary context. Intuitive Comprehension has to do with the truth-conditional knowledge that a hearer can achieve about an utterance independently of the information that pertains to the specific context of that utterance.

I then show that, even when it is impossible for them to carry out any pragmatic processes in specific contexts, speakers can nevertheless achieve at least an *intuitive* comprehension of utterances of semantically under-determined sentences. This is apparently accomplished in virtue of the fact that hearers can expand the minimal semantic content delivered by formal semantics alone by employing some very general constraints, which may be described as semantic/conceptual in nature. For example, in figuring out the intuitive truth-conditions of an utterance of "It's raining", hearers may expand the minimal content yielded by formal semantics alone, *<that it's raining>*, by employing the semantic/conceptual constraint to the effect that every episode of rain-falling necessarily happens at some location, and thus derive that that utterance is true (in any context of utterance) iff it's raining at some location; in figuring out the intuitive content of "Jill is ready", they may apply to the minimal, formal semantics derived content *<that Jill is ready>* the semantic/conceptual constraint to the effect that every instance of readiness is necessarily readiness for some purpose, and thus derive that that utterance is true (in any context) iff Jill is ready for some purpose.

The constraints I have mentioned are both conceptual and semantic, in that they constrain not only concept-use, but also the correct use of terms like "rain", "ready", "green" in ordinary linguistic practices. Nevertheless, I will generally prefer to address them only as *conceptual constraints.* That is because I believe that the conceptual domain is more stable than the semantic domain, insofar as one could be competent about a concept *F* by at the same time having false beliefs on the meaning of the term *e* which is associated with *F* in a certain language *L*; by contrast, if one is competent on the meaning of a certain term *e* in language *L*, one is generally also competent on the concept *F* associated with this meaning. In other words, the conceptual domain is cross-linguistic and more stable than the semantic domain, which is tied to linguistic competence. In most of the cases I will scrutinize, conceptual constraints and semantic constraints will coincide *de facto,* for I am acting on the assumption that speakers and listeners are competent users of the English language. For them, knowing the meaning of "rain" entails that they are also competent about the concept RAIN.

So, conceptual constraints coincide with semantic constraints for every competent user of a certain

_____

hearer listens to the sentence "I've had breakfast" and has no access to a context in which pragmatic processes could be carried out, then she still can understand the sentence for what it strictly speaking means, namely *<that the speaker has had breakfast at least once in her life>*. Thus, failure to run pragmatic processes does not prevent comprehension when the under-specificity at issue is "shallow" or merely pragmatic.

expression *e* in language *L*. The peculiarity of this kind of constraints is, however, that they operate on an expression's semantics without being necessarily connected with the *syntax* of the expression. I will argue that the application of conceptual (semantic) constraints doesn't respond to syntax, but it rather responds to pre-determined "schemata", or ways of organising thought and discourse about matters such as rain-falling, being ready, etc. Cognitive schemata in particular determine what I call *standard cognitive formatting*, i.e. standard structures for thought.

The following image illustrates the way I conceive the relation between syntax, semantics, conceptual constraints and pragmatics. Syntax concerns the relations between the terms of a language *L*, and determines which expressions are well-formed and which are not. Semantics concerns the truth-conditions of utterances of (well-formed) sentences in *L*. Conceptual constraints concern the correct way of structuring thought about any matter. Conceptual constraints *coincide* with semantic constraints for every speaker who masters the meaning of an expression *e* in language *L*. They, however, have a wider scope, in that they may be involved in other cognitive activities which do not directly relate to the interpretation of utterances. Finally, pragmatics concerns the use of language and supplies the means to determine which utterances enjoy distinctively pragmatic virtues such as informativeness, relevance, perspicuity, clarity, etc.

Syntax	Semantics	Conceptual constraints	Pragmatics
Relations between expressions of *L*.	Truth-conditions of utterances.	Organisation of thought.	Use-related virtues of utterances.
	Semantic constraints = Conceptual constraints for competent users of any expression *e* in *L*.		

Let us now go back to the notion of intuitive comprehension. My hypothesis is that intuitive comprehension reveals the hearer's capacity to reconstruct the content and truth-conditions of an under-determined utterance by employing the pre-determined "scaffolding" provided by the standard cognitive format that regulates thoughts (and discourse) on the relevant matters: for example, the standard cognitive format for thought and discourse about rain-events, about states of readiness, of greenness, etc. I conclude by stressing the fact that conceptual (semantic) constraints might have a fundamental role in the comprehension of semantically under-determined sentences, even when this comprehension is about utterances in specific contexts to which the hearer has access. In this perspective, the Radical Contextualist Response, which is formulated solely in terms of free pragmatic processes, looses much of its appeal in favour of a response to the Comprehension Problem that explicitly envisages the contribution of *constraining elements that are conceptual (semantic) in nature*. At the end of **Chapter 2**, I therefore endorse the view that conceptual (semantic) constraints play a role in the comprehension of utterances of under-determined sentences (even in specific contexts):

[The Conceptual Constraints Claim] Conceptual (semantic) constraints play a role in the comprehension of utterances of under-articulated sentences.

The claim is still very much provisional and needs further qualification. In particular, it immediately elicits a question as to *the status* of these constraints. Conceptual constraints may be characterised as ways of organising information (and hence also thought and discourse) about the world. A claim based on this characterisation may however sound too simplistic and commonsensical. All humans organise information about the world according to certain constraints; however, the genuine question is *how* and *at which level* these constraints affect utterance comprehension.

The issue is challenging, insofar as one may observe that talk of conceptual (semantic) constraints could be dispensed with, in favour of talk about *syntactic* constraints. Appealing to syntactic

constraints may seem to guarantee a more "compact" explanation of how utterance comprehension is possible in cases of under-articulation, for syntactic constraints undoubtedly fall under the domain of linguistic competence. Several authors in the last decade have put a strong emphasis on the alleged role of syntactic constraints in the comprehension of semantically under-determined sentences like "It's raining", "Jill is ready", "The leaves are green" etc. (see Stanley 2000, Stanley & Szabó 2000, Szabó 2001, King & Stanley 2005). I shall therefore explore the possibility that what I have introduced as conceptual (semantic) constraints are really *syntactic* constraints on interpretation. As it will emerge, the proposal has more than one weakness.

In **Chapter 3**, my focus shifts to a less radical brand of Contextualism, known as Indexical Contextualism. The focal claim of Indexical Contextualism is that all the processes by which an utterance's content is supplemented contextually are traceable to the sentence's logical form (see Stanley 2000). Since there is no process of content supplementation that is not controlled at the syntactic level, all such processes are reduced to forms of indexicality resolution broadly conceived, i.e. of satisfaction of empty positions at some level of syntax. So for example, the process by which the content of an utterance of "It's raining" is supplemented with a locational element like *<Paris>* is supposed to be triggered by the presence of a *hidden argument* for locations in the logical form of the predicate "to rain". The predicate is postulated as denoting a relation between times and locations *rain(t, l)*, where the locational argument is not articulated at the superficial grammatical level but nevertheless marks a position in deep syntax. With this picture in the background, we can suppose that the Indexical Contextualist will address the Comprehension Problem in the following way:

[The Indexical Contextualist Response] The comprehension of utterances of semantically under-determined sentences necessarily depends on the resolution of covert forms of indexicality (e.g. saturation of hidden slots in logical form).

Before moving on, let me make a clarificatory remark: I already argued in **Chapter 1** that semantic under-determinacy should not be conflated with context-dependency and, specifically, with indexicality. In addressing the Indexical Contextualist Response, my purpose is certainly that of indirectly supporting this claim. That is, in demoting the claim that conceptual constraints are reducible to mere positions in logical form, I aim at corroborating the point that semantic under-determinacy is not a species of indexicality.

Let us then return to the Indexical Contextualist Response. In order for the Indexical Contextualist Response to be successful, it has to be admitted that sentences such as "It's raining", "Jill is ready", "The leaves are green", which we have characterised as ultimately under-articulated, all really share a covertly indexical structure, which means that the predicates that occur in them feature some extra-argument place at a deep level of syntax. Thus, the predicate "to rain" has to have the structure *rain(t, l)*, which includes a hidden argument for locations *l*; the predicate "being ready" has to have the structure *ready(x, y)*, where *y* is a purpose- or activity- argument; the predicate "being green" has to have the structure *green(x, y)*, where y is a respect- or part- argument. This all has to be accepted in order for the Indexical Contextualist response to get off the ground. However, are there any good reasons to accept all this? Are there any good reasons to revise our views concerning the syntactic structure of predicates like "to rain", "being ready", "being green"? The aim of **Chapter 3** is precisely that of questioning the reasons that proponents of Indexical Contextualism advocate in support of this "revisionary" proposal.

The most interesting defence of Indexical Contextualism is known as the Binding Argument (see Stanley 2000). Very sketchily, the Binding Argument purports to establish the Indexicalist position about the structure of sentences like "It's raining", "Jill is ready" etc., by means of the following reasoning. First, a sentence like "It's raining" can be embedded into a more complex sentence which contains a quantification over locations, as for instance "Everywhere I go, it rains". This complex

sentence gives rise to a bound reading, in which a putative locational variable is controlled by the quantifier "everywhere". Under this reading, the sentence expresses the content *<for every location x where the speaker goes, it's raining in x>*. The crucial step in the argument then consists in saying that, if the bound reading is possible, then there must be a bindable variable articulated in the logical form of the embedded sentence "... it rains". If the bindable variable is present in the embedded sentence "... it rains", then it must be present in the sentence even when it occurs unembedded. In conclusion, the sentence "It rains"/"It's raining" must contain a hidden indexical element corresponding to a locational variable.

This argument has been criticised on many fronts. My criticism could be thus summarised. The Binding Argument allegedly shows that the syntactic structure of predicates such as "to rain" has more complexity than meets the eye, in that it contains a hidden extra-position. This assumption at the syntactic level, however, brings to an incorrect account of all those occurrences of the predicate "to rain" in which the predicate clearly is used with a non-complex, monadic syntactic structure: call these cases "Zero Readings". Consider, for example, the following sentence:

(27) Every time water falls from the sky, it rains.

In this sentence, the predicate "to rain" seems to have the simple, monadic structure *rain(t)* where raining is a property of times, or events; that is, the predicate seems to involve no hidden locational argument place. Indeed, (27) could be read as expressing the content *<every time* t *water falls from the sky, a raining-episode occurs at* t>. If, however, the predicate did contain a hidden locational argument thus having a dyadic structure *rain(x, y)*, then (27) would forcedly become an open sentence, whose content is *<every time* t *water falls from the sky, it rains in* l *at* t>. *This* content is unevaluable unless a value is provided for the free variable *l*. However, it seems that (27) is perfectly evaluable with no need for any value assignment. The Indexicalist could respond to this challenge by interpreting (27) as having a bound reading. For instance, by supposing that the quantifier "Every time water falls from the sky" binds a location variable in the embedded clause "it rains", as in *<For every time* t *and every place* l *in which water falls from the sky, it rains at* l *at* t>. This reading seems fine, however, two considerations are called for: firstly, the bound reading is available even for the theorist who doesn't believe that "to rain" has structure *rain(x, y)* (for instance an event-semanticist: see below for further details); secondly, what really should trouble the Indexicalist is the fact that zero-readings cannot be explained with the idea that what *seems* a monadic predicate *is* indeed a monadic predicate *rain(t)*. The Indexicalist can't have a direct explanation of zero-readings—simply in terms of monadic predicates. Rather, he has to resort to bound readings in order to comply with the intuition that sentences like (27) are not open sentences. By contrast, the non-Indexicalist can easily account for zero-readings by simply pointing at the fact that a predicate like "to rain" has a monadic syntactic structure *rain(t)*.

These considerations lead to the conclusion that postulating a complex syntactic structure for predicates such as "to rain" (but also "being ready", "being green", etc.) on the account that this allows to explain bound readings blocks the possibility of adequately explaining *zero-readings*. Adopting an event-analysis and semantics like the one suggested by Davidson (1967), Parsons (1990), Cappelen and Hawthorne (2007) and Recanati (2010) could probably guarantee an acceptable account of *both* bound readings and zero-readings. I therefore sympathise with all those authors who give up any assumptions as to the hidden syntactic structure of predicates like "to rain" and prefer to see bound readings (and zero-readings) as displaying relations between structural aspects of events (for instance, between their spatial and temporal aspects). This, however, is not the main point of the chapter. The crucial point is that, if an event-analysis is taken on board, then it becomes apparent that the Indexical Contextualist's commitments as to the syntactic form of expressions are dropped in favour of a commitment to aspects that reflect the way we organise information about events of raining, of being ready, of being green *at a genuinely*

conceptual/semantic level. The fact that the internal structure of events of raining, but also supposedly of states of readiness and of states of greenness enters into the way we make sense of the semantics of sentences (in both bound readings and zero-readings) may reveal an interaction between the way we organise thought and discourse about the world and the way we interpret sentences. Ultimately then, *abandoning syntactic commitments may bring the theorist to embrace commitments at a conceptual/semantic level,* thus going in the direction of an account which regards conceptual constraints as positively involved in the comprehension of under-articulated sentences, at the expense of syntactic constraints.

In **Chapter 4**, I consider the idea that conceptual constraints may be reduced, at a semantic level, to parameters for evaluation. This idea implies essentially engaging with forms of Truth-Relativism. Semanticists have so far identified two forms of Truth Relativism—a "moderate" and a "radical" form. Both the "radical" and the "moderate" versions of Truth Relativism postulate that the truth-value of a sentence is relative to a circumstance of evaluation that encompasses more parameters than just a possible worlds- parameter: this common formal move I call the "Extra-Parameters Move". Thus for example, Temporalists (Prior 1968, Kaplan 1989) postulate that tensed sentences are true with respect to a possible world-parameter w and a time-parameter t; Relativists about epistemic modality (Egan, Hawthorne, Weatherson 2005, Egan 2007, Egan 2011) have it that sentences of the form "It might be that φ" are true relative to a possible world-parameter w and an epistemic standard-parameter e.

The aspect on which moderate and radical relativists diverge has to do with *how the value of the extra-parameters is fixed*. Moderate Relativists, also known as Non-Indexical Contextualists (Kaplan 1989, Lewis 1980, Kölbel 2002, Recanati 2007) have it that the value of extra-parameters is fixed at the circumstances of the context of utterance; Radical Relativists have it that the value of the extra-parameters is fixed at the circumstances of a context of assessment, which may be different and completely independent of the context of utterance (see MacFarlane 2005a).

As I mentioned in **section 2** as a merely "sociological fact", theorists have explicitly endorsed or at least foreshadowed a Truth Relativist approach as a way of opposing Indexical Contextualism. A potential route for those who are unconvinced by the conceptual constraints view and by Indexical Contextualism would therefore be that of trying to *reduce conceptual constraints to parameters of evaluation*.

Supposing that conceptual constraints are parameters of evaluation entails endorsing a view about the truth-conditions of under-determined sentences which is compatible with something like the following statements: An utterance of "It's raining" is true relative to a possible world-parameter w and a location-parameter l; an utterance of "Jill is ready" is true relative to a possible world-parameter w and a purpose-parameter p; an utterance of "The leaves are green" is true relative to a possible world-parameter w and a respect-parameter r. I presume that the Truth Relativist would therefore address the Comprehension Problem by giving the following response:

[The Relativist Response] Comprehension of utterances of semantically under-determined sentences necessarily requires recognising a form of truth relativity.

I approach critically the Relativist Response by assessing whether this kind of solution is *enough motivated* for the purposes of capturing conceptual (semantic) constraints. I first consider Radical Relativism and subsequently Moderate Relativism, also known as Non-Indexical Contextualism. Concerning Radical Relativism, I survey the two main motivations that have led theorists to endorse this view in recent years: (1) capturing disagreement on matters "of inclination" and (2) capturing retraction on matters "of inclination" (see MacFarlane 2005a, 2007b, MS.). I argue that neither of these motivations has any bearing on the issue of capturing conceptual (semantic) constraints. Radical Relativism is not the best option for accounting for these constraints, because the notion itself of "assessment-sensitive" truth is tightly linked to the "perspectivality" of truth i.e. to the

dependence of truth on points of view that cannot be reduced to any objective aspect of the world, such as taste standards, epistemic states-standards, moral-code standards. However, conceptual (semantic) constraints do not *in general* concern perspectives; they may at best be regarded as determining a truth-predicate which is merely "feature-sensitive", i.e. which is relative to objective aspects of the world, such as times, locations, comparison classes, purposes, parts, respects etc. The best option to account for conceptual (semantic) constraints within an extra-parameters framework may therefore be a Moderate Relativism—also known as Non-Indexical Contextualism.

The problem with Non-Indexical Contextualism, though, is that it gives a strained account of what semantic under-determinacy is. In particular, Non-Indexical Contextualism forces the theorist to radically re-frame the intuitions of semantic under-determinacy with which all theorists in this area of investigation start. More to the point, if Non-Indexical Contextualism is true, intuitions as to the semantic under-determinacy of sentences like "It's raining", "Jill is ready" etc. cannot be intuitions about content, but they have to be intuitions about the utterance's circumstances of evaluation. This radical re-framing is, however, methodologically suspect: it doesn't seem fair that a theory which starts with a certain pool of *data* winds up distorting those initial data as a result of the theory itself having been established. Thus, it seems that we have some reasons to reject Non-Indexical Contextualism too, as an account of conceptual (semantic) constraints.

Given the failure of both Radical Relativism and Non-Indexical Contextualism to capture conceptual (semantic) constraints, and given also the previous failure of Indexical Contextualism, I conclude that we should abandon any attempt to reduce such constraints to either syntactic components, such as hidden positions in deep syntax, or to aspects of evaluation, such as extra-parameters for the relativisation of truth-value. Rather, we should take conceptual (semantic) constraints as operating at an independent level. In **Chapter 5**, I will provide a picture in which conceptual (semantic) constraints operate on an utterance's truth-conditions not because they obey to any syntactic restrictions or restrictions about evaluative aspects, but because they obey to restrictions imposed by "schemata" or frames, which are ways of organising thought and discourse about the world.

5. The Conceptual Constraints View

In the preceding sections, I suggested that we should move away from a reduction of conceptual constraints either to aspects of syntax or to parameters for evaluation. Instead we should be ready to approach these constraints as operating on an independent level.

I take conceptual constraints to be primarily *ways of structuring and organising thought about certain matters.* We structure thought about metereological matters in terms of phenomena such as rain-falling, snow-falling, etc. that concern locations; we structure thought about being ready in terms of purposes or activities for which one is ready; we structure thought about being green in terms of being green in some respect, or part, or against some comparison class. That thought about these matters is structured thus may be explained by the fact that subjects belonging to a certain (environmental, cultural, social, linguistic) group need to think in the same terms in order for their thought to adequately fit reality and the way other subjects think about and describe this reality. In figurative terms, subjects that belong to the same (environmental, cultural, social, linguistic) group need to share a *standard cognitive format* for thought about certain matters. To characterise the phenomenon in terms familiar to cognitive psychologists, we may view conceptual constraints as evoking very general and fundamental *schemata*[6], i.e. structured configurations divided into different "dimensions" or "slots" that have to be completed with specific kinds of values. The attainment of the appropriate cognitive formatting could be depicted as a mechanism of cognitive "slot-filling".

In this thesis, I will be mainly talking about *conceptual* constraints, rather than semantic constraints. The reason for this choice is that I believe that the conceptual domain is *more stable* than the

---

[6] For the notion of *schema*, see Rumelhart and Ortony 1977, Rumelhart 1980, Cohen and Murphy 1984.

semantic domain. For example, it could be the case that one is competent on a concept *F* but at the same time one has false beliefs about the meaning of a word *e* which, in a language *L*, is associated with *F*. Or, it could be the case that one who doesn't master the meaning of a certain word *e* in a language *L* (which, by hypothesis, is associated with concept *F*) manages to comprehend that word with the help of her mastery of precisely *F*. The upshot of these considerations is that conceptual constraints are not as tied to language and linguistic competence as semantic constraints are: being cross-lingustic, they seem more stable and reliable.

It is however important to keep in mind that, when one is a competent user of a word *e* in language *L*, one certainly also masters the concept *F* associated with *e*. *So for competent users of any expression* e *of* L*, conceptual constraints and semantic constraints coincide, as a matter of fact*. For instance, knowing the meaning of the word "rain" in English entails mastering the concept RAIN associated with it. In all the cases I will consider ("It's raining", "The leaves are green", "Jill is ready", etc.), it will be a tacit assumption that the hearers and speakers involved are competent users of the English language and, as a result, also competent users of the relevant concepts. Thus, even though in what follows I will tend to use the the notion of conceptual constraints at the expense of that of semantic constraints, I will nevertheless maintain that, in the specific cases I will scrutinise, talk of conceptual constraints might be safely substituted by talk of semantic constraints, since I will assume that hearers (and speakers) are competent language users.

Let us now turn to the question of how conceptual constraints affect the comprehension of utterances of semantically under-determined expressions. On the assumption that hearers master the relevant constraints, we can suppose that *hearers exploit conceptual constraints* (together with information about the context, the speaker's intentions and so on) in order to reconstruct the content and truth conditions of utterances of under-determined sentences. Let us see how this might happen by making an example.

Suppose a speaker utters "Jill is ready". This sentence is semantically under-determined. Provided that the hearer is a competent user of the predicate "being ready" as well as of the concept READY, she can use the conceptual (but, in this case, also semantic) constraint that readiness is readiness for some purpose/activity, plus information that she may have gathered from the context concerning Jill's particular purpose, in order to reconstruct the truth-conditions of this utterance. The hearer may be in a position to perform this reconstructive task insofar as these conceptual constraints are part of the cognitive format for thought and discourse that she (the hearer) and the speaker presumably share. A precondition for doing this is certainly that the hearer has at least the background presupposition that the speaker is a competent language- and concept-user, i.e. that she structures thought and discourse about being ready in terms of being ready for some purpose. If this presupposition is in place and there are no reasons for defeating it, then the hearer can proceed applying her own conceptual/semantic competence.

The interpretation of an utterance of "Jill is ready" may therefore be regarded as resulting from a combination of competences: first, there is a subject's mastery of a compositional, syntax-driven semantic system, which plausibly constitutes what we call "linguistic competence"; secondly, there is one's conceptual/semantic competence, which integrates the results of linguistic competence; third, there is one's capacity to gain information from the context, which may provide additional propositional material. So for instance, the hearer may employ her linguistic competence in order to decode the sentence's content *<that Jill is ready>*; by exploiting the conceptual (semantic) constraint that readiness is such only with respect to a purpose/activity, she may "expand" such a minimal content so as to obtain a suitable extra conceptual/semantic slot. Context may provide her with information concerning the value with which the slot should be filled—for example, skiing. The combination of conceptual/semantic constraints application and contextual information allows her to derive the content *<that Jill is ready for skiing>*.

So far I have described the way the process may be taken to unfold, but I haven't said anything concerning its nature. Is this process semantic or pragmatic? Is it inferential or non-inferential? Is it

conscious? Does it involve any particularly sophisticated cognitive abilities?

"Conceptually constrained" processes depart from semantic processes in that they do not exclusively govern the interpretation of linguistic constructions, but are supposedly involved in enterprises that are not purely language-related. Nevertheless, conceptual constraints *are* semantic constraints in a very important sense. As I have already made clear, conceptual constraints *coincide* with semantic constraints for competent users of any expression *e* of language *L*. If this is the case, then the processes triggered by such constraints doubtless belong to semantics. The peculiarity about these constraints is that, even though they determine genuinely semantic, truth-conditional effects, thus helping to figure out "what is said" by an utterance, they do so in a way that is *unconstrained by syntax*. Thus for instance, the semantic content of "It's raining" established compositionally on the basis of syntax alone is the proposition *<that it's raining>*. The application of conceptual constraints "expands" the conceptual/semantic structure of that content so as to obtain something like *<that it's raining somewhere>*. Note however that this process of "expansion" unfolds independently of any constraints posed by the syntax of the sentence. Conceptual/semantic constraints therefore determine *semantic, but syntactically unconstrained processes*. The following chart illustrates the way I conceive the relation between semantics, syntactic constraints and conceptual/semantic constraints on the one hand, and pragmatics on the other. As one can see, both syntactic constraints and conceptual constraints (together with contextual information, where appropriate) concur to determine the truth-conditions of an utterance, or "what is said" by it: this licenses the claim that they both give a contribution to the semantics of the utterance. Pragmatics, on the other hand, concerns what the utterance conveys rather than say.

What is said (Semantics)	Syntactic constraints	"It's raining" is true iff it's raining.
	Conceptual (semantic) constraints	"It's raining" is true iff it's raining *somewhere* (contextual info may be plugged in, if appropriate: e.g. *in Paris*);
What is implied (Pragmatics)	Conversational principles/maxims	Any implicated propositions.

More on the nature of conceptually constrained processes will be said in **section 7** of **Chapter 5**. Conceptually constrained processes are to be viewed as inferential in the standard sense of the term, i.e. as involving a set of premises that lead to a conclusion *via* a valid pattern of inference. The performance of the inference is to be understood as implicit, or tacit. This inferential conception is to be contrasted with views of processes of utterance comprehension as purely translational (see Millikan 2004), as associative (see Recanati 2002b), or as guided solely by cognitive mechanisms such as "path-of-least-effort" mechanisms (see Sperber and Wilson 1986/1995). I refer to **section 3** of **Chapter 5** for these issues.

Moreover, even though conceptually constrained processes are inferences which hearers run in order to work out the truth-conditions of an utterance in a certain context, they need not be inferences that result in *attributions of attitudes* to the speakers, where typical attributions of attitude are of the form "X means that P", "X believes that P", "X intends that P". For example, the conceptually constrained interpretation of Ann's utterance of "Jill is ready" needs not produce as its result the proposition *<that Ann means that Jill is ready to ski>*. If the conceptual constraints view is correct, the hearer who interprets an utterance of "Jill is ready" only needs to go through a "mechanical" process of expansion of the sentence's conceptual/semantic structure, so as to obtain an extra conceptual/semantic slot which she may fill-in with information gained from the context of utterance. This "mechanical" process of structure-expansion and slot-saturation doesn't need to involve the subject's meta-representational capacities. Granted, the hearer may exercise her meta-representational capacity to enquire some background aspects of the situation of utterance—for

example, to discover what the speaker is talking about, what she's making reference to, and so on. This, however, does not entail that the result of the *overall* utterance interpretation need be an attribution of attitude or intention of the form *X means/believes/intends that p*. In other words, even if meta-representation may still be conceived as part of the interpretive process, this fact needs not affect the structure of the interpretive process conclusion, which may remain an object-level representation. The view of conceptually constrained processes as inferences having object-level conclusions contrasts with the Gricean model, which envisages interpretive inferences that are meta-representational through and through, and also with the Relevance-Theoretic model (see Sperber and Wilson 1986/1995, Carston 2002), which envisages inferences whose premises are not necessarily meta-representational, but whose conclusion is. More will be said on this issue in **section 4** of **Chapter 5**.

Further developments for the conceptual constraints view will presumably have to ascertain how far a conceptual constraints account can be extended: do we need conceptual constraints just to explain the comprehension of semantically under-determined sentences? Or do conceptual constraints play a role also in the reference-assignment of pronouns, in the interpretation of compound expressions, idiomatic expressions, figurative speech and so on? Another interesting topic of investigation concerns the extent to which users that are, by hypothesis, employing conceptual constraints produce meta-representations as the results of their comprehension process. I reserve these challenging developments for future work.

Chapter 1

On Semantic Under-determinacy

1. Preliminaries

In the present work, I will be concerned with the semantics of sentences like the following:

(1) It's raining;
(2) Jill is ready;
(3) Lisa is tall;
(4) The leaves are green;

What is interesting about these examples? As I already anticipated in the **Introduction**, these sentences contain expressions which suffer from a form of semantic under-specification, and yet are not clear examples of, for instance, ambiguity, vagueness, or context-dependence. For convenience, I have decided to call these expressions *semantically under-determined*. In this section, my purpose is that of fostering some pre-theoretic intuitions about their defectiveness.

Before proceeding, it will be appropriate to fix some terminology. Following the standard contemporary conception—which traces back to Kaplan (1989), Lewis (1980) and Stalnaker (1970) —I will call *semantics* the formal device whose task is to determine the truth-conditions of a sentence in context, at (at least) a possible world. Semantics standardly takes as its input a couple formed by a sentence *s* and a context *c*, and delivers the truth-conditions of *s* at the world of the context w_c. This process follows a principle of *compositionality*, according to which the meaning of a sentence in context is determined by the meanings (in context) of its components and their syntactic arrangement.

In this framework, I will talk of *semantic content* as the output of these compositional processes, which is standardly to be evaluated at a possible world. I will use as equivalents of "semantic content" also expressions like "sentence meaning", "what is strictly speaking said", "the semantics of the sentence", etc. Semantic content is determined solely by the content (in context, where appropriate) of the sentence's components, according to the way these components are syntactically arranged. *Syntax* provides the rules by means of which well-formed sentences can be set apart from ill-formed ones, and acts as a constraint on interpretation in the sense that *only* what is syntactically articulated in a sentence goes into the sentence's interpretation. So for example, the semantic content of (1) results from composing together the meanings of *It, is, raining*, thus resulting in the content *that it's raining*. This content is what should be evaluated at a circumstance—usually a possible world, notably the world in which the utterance of (1) is performed. Similarly, the semantic content of (2), resulting from composition of the words' meanings, is *that Jill is ready*. The semantic content of (3) is *that Lisa is tall*, while that of (4) is *that the leaves are green*.

Semantic content thus conceived—as the result of compositionality-driven interpretation—often doesn't deliver any propositional item, where a *proposition* is a content (in context) which is amenable to evaluation in terms of truth and falsity (at at least a world). Whenever the semantic content of a sentence *s* doesn't deliver an evaluable content (in context) for utterances of *s*, and (as already mentioned in the **Introduction**) this apparent semantic defectiveness is not *prima facie* imputable to any explicit form of ambiguity, vagueness, indexicality or ellipsis, I shall say that

semantic content gives rise to semantic under-determinacy. In this section, I will confine myself to providing some examples of the phenomenon, to make it vivid to the reader.

Imagine that, at time *t*, you happen to hear an utterance of (1). The semantics of English enables you to know that what has been said is *that it is raining at t* (where the specific time *t* is plausibly incorporated into the present tense). In most cases, upon hearing an utterance of (1), it would take just a look outside the window to check whether it is true. However, suppose you looked out of the window but you saw no rain falling. It doesn't follow that the speaker's utterance of (1) is false. It could be raining somewhere else, e.g. in Paris, and the speaker may be talking about precisely that place. But how would you know that? No specifics concerning the location of the rain are provided by what (1) strictly speaking says, i.e. by its semantic content. In this case, one shall say that the meaning of (1) doesn't say how the world should be in order for an utterance of it to be true.

Suppose now you hear an utterance of (2). Unless you are sufficiently informed about what the speaker is talking about, you are probably going to find this utterance poorly informative; most prominently, you are likely to feel at a loss as to whether what has been said is true or false. What the sentence means, i.e. the semantic content *that Jill is ready* fails to determine how the world should be in order for an utterance of (2) to be true, because it fails to specify *what activity* Jill is ready for.

Now consider an utterance of (3) made by a speaker who is referring to Lisa, a 25-year-old woman. It should not be difficult for you to understand what the speaker says, nor to know whether the utterance is true or false: you only need to know who Lisa is, and what her approximate height is. Indeed, language users do not have any difficulty in understanding and evaluating utterances like that of (3): but that, I submit, is because there is a *default standard* for being tall, which is the average height of humans of the community one belongs to. Suppose for a moment this default standard was not operative: the speaker is talking about Lisa's height for a basketball player. Now the situation changes. For, even if Lisa is a tall person, she may fail to be tall for a basketball player. If this were the case, the speaker's utterance of (3) would be false. So, is it so straightforward in the end that Lisa is tall? What standard has to be taken into consideration in order for this judgement to be true? Again, it seems that the meaning of (3) (*that Lisa is tall*) cannot be of help. No salient standard is specified at the level of semantic, compositionally-driven interpretation. So it seems that the semantic content of (3), just like in the previous cases, doesn't fully specify how the world should be in order for an utterance of (3) to be true.

Finally, imagine that you hear an utterance of (4), while the speaker is referring to the leaves of the plant which is sitting in the corner of her living room. You may be able to evaluate (4) just by seeing what the colour of the leaves is. In general, language users establish the truth or falsity of sentences of the form "x is green" by looking at how x looks. However, suppose the plant in question looked green on the outside, but just because a child painted it green; under the paint, the plant's leaves are actually red. It seems now that the utterance of (4) is false: strictly speaking, the leaves are not green, even though they look green because they were so painted. In light of these facts, is it now so obvious that the leaves are green? In *what way* should the leaves be green in order for (4) to say something true? The meaning of (4) (*that the leaves are green*) cannot specify in which sense the leaves have that colour: further details from the context of utterance are required. So again, and just like in the previous examples, the semantics of (4) seems to fail to specify how the world should be in order for an utterance of (4) to be true.

What examples (1)-(4) show is that some of the sentences we use seem to fail to specify the truth-conditions of their utterances. The idea is that, if one is to hear an utterance of any of these sentences without being provided with further contextual information, one is left in the dark as to what must be the case in the world in order for what one has heard to be true. What sentences say in virtue of their meaning alone is "too poor" in order for utterances of those sentences to be evaluated. In the following section, I will be concerned with introducing the reader to the historical framework in which the idea of semantic under-determinacy has risen.

2. Historical background

2.1 Grice and the Saying-Implicating distinction

Grice's seminal work on the semantics of logical connectives such as "and" and "or" is often cited as responsible for having introduced a very sharp separation between what is conventionally "said" by a sentence and what is communicated, or "implicated" by an utterance of that sentence. The Gricean distinction has it that what a certain sentence says is the result of the combination of the words' conventional meanings. In *Logic and Conversation*, Grice writes: "I intend what someone has said to be closely related to the conventional meaning of the word (the sentence) he has uttered." (Grice 1967/1989: 25). In addition to the conventional element, what is said is strictly bound to the elements explicitly articulated in the sentence's syntax. Grice makes it clear that what composes what is said must track the way the sentence is arranged, in that it must correspond to "the elements of [the sentence], their order, and their syntactic character" (Grice 1969/1989, p. 87). Note that, even if a sentence's meaning is conventional, this does not imply that it does not take the context of use into account. It is compatible with the Gricean picture that the semantics of some words, like for instance indexicals such as "I" and "you" or demonstratives like "he", involves a form of contextual saturation in order for their reference to be fixed; ambiguous terms as well require contextual disambiguation. Here is a passage in which it is clear that the semantic value of both such kinds of terms fully belongs to what is conventionally "said":

> Suppose someone to have uttered the sentence "He is in the grip of a vice" […] [F]or a full identification of what the speaker had said, one would need to know: (a) the identity of x; (b) the time of the utterance; (c) the meaning, on the particular occasion of utterance, of the phrase "in the grip of a vice". (Grice 1967/1989: 25)

On point (a), Grice seems to grant that, in order to understand what has been said, one needs to assign a referent (in context) to the occurrence of "he"; while on point (c), Grice seems to be admitting that the meaning of the phrase "in the grip of a vice", disambiguated in context, is to become part and parcel of what has been said as well. So, that context is taken into account does not mean that linguistic conventions are being less applied in figuring out what is being said. In general, then, we can say that what a sentence says is the result of applying a set of linguistic conventions, some of which do not dictate to take the context into consideration, some of which do.

Besides the notion of what is said, Grice leaves room in his framework for contents that can be communicated without being said. Such further contents are called by Grice *implicatures*. An example of implicated content may be the following: A and B are visiting a friend. After a a couple of hours, A announces "It's 6 o'clock" to B. B can interpret A as saying *<that it's 6 o'clock>* in order to suggest or *implicate* that, for instance, it's time to leave. As it can be noticed, an implicature is a proposition inferred from the fact that a certain sentence ("It's 6 o'clock") with a precise conventional meaning (*<that it's 6 o'clock>*), has been uttered. What is communicated via implicature is a distinct content from that expressed by sentence meaning alone. This additional content normally has no bearing on the truth-condition of what is conventionally said by the sentence: it is, in general, truth-conditionally distinct and autonomous.

How are implicatures generated from conventional meaning? In general, Grice maintains that an implicature is built up from an already established sentence meaning *plus* some contextual pieces of information. However, he admits that some implicatures are directly *built into* the very meaning of a term; that is, nothing more than the conventional meaning of the term employed is needed in order to generate the implicature. This is what he calls a *conventional* implicature. An example may be provided by the word "but". "But" gives the same truth-conditional contribution as "and": hence,

what a sentence like "He is rich, but generous" can be regarded as saying is that the demonstrated person is rich and generous. Nonetheless, the mere use of a word like "but" commits the speaker to (and prompts the hearer to understand) the proposition <*that B is rich and*, in contrast with this, *generous*>. Since this proposition is not strictly speaking said, Grice concludes that it is implicated by convention, i.e. by the mere fact that the expression "but" is used.

Other kinds of implicature, not directly generated by the semantic properties of expressions, arise in connection with the idea that having a conversation is a matter of cooperation. A conversation is usually not an exchange of linguistic expressions totally disconnected from each other. At any stage of the conversation, there shall be some expected, appropriate contributions, that speakers should give. In normal conditions, speakers will therefore try to make it that their contribution to the overall conversational enterprise is as appropriate as possible; hearers too will struggle in order to interpret the moves that their interlocutors make as most appropriate and helpful for the sake of the conversation. In general then, it seems that being engaged in a conversation involves what Grice calls the Cooperative Principle:

> Cooperative Principle: Make your conversational contribution such as is required, at
> the stage at which it occurs, by the accepted purpose or direction of the talk exchange
> in which you are engaged. (Grice 1967/1989: 26).

There is more than one aspect in which a conversational participant's contribution might be appropriate—or not appropriate. Grice individuates four of these aspects, which he calls "categories" (borrowing the term from Kant); to each of these categories there corresponds a maxim (accompanied by a set of sub-maxims) for the related appropriate behaviour in a conversational exchange. The first of these four maxims pertains to the category of Quantity, i.e. the amount of information that one is allowed and encouraged to provide in the course of a conversation:

> Quantity: (i) Make your contribution as informative as is required (for the current
> purposes of the exchange). (ii) Do not make your contribution more informative
> than is required. (*ibid*: 26);

The second maxim is about the Quality of the contribution that one is allowed to provide. Here it is assumed that the quality of a true utterance is better than the quality of a false one, and that the quality of a warranted or justified utterance is better than the quality of an unwarranted one.

> Quality: (i) Do not say what you believe to be false. (ii) Do not say that for which
> you lack adequate evidence.

The third maxim concerns the Relation of one's contribution, presumably with respect to the previous contributions and to the general goal of the conversation: The idea is that one's contribution must relate in some clear way with these two elements, hence the requirement of relevance.

> Relation: (i) Be relevant;

Finally, the fourth maxim covers the category of Manner, that is, not the category of "what" is said but that of "how" what is said is said. Here the super-maxim, comprehensive of the other more specific sub-maxims, is "Be perspicuous". The sub-maxims are:

> Manner: (i) Avoid obscurity of expression; (ii) Avoid ambiguity; (iii) Be brief
> (avoid unnecessary prolixity); (iv) Be orderly.

How is the notion of implicature connected to the maxims? In general, an implicature may arise every time a speaker, by saying what she says, fails to fulfil one of the maxims, even though there is no reason for the hearer to suppose that she is not being cooperative. There are at least two ways in which a speaker might fail to satisfy a maxim. First, the speaker may commit an "unnoticeable" or "quiet" violation, such as that in the following dialogue:

(5) a. Alice: "Shall we go to the movies tonight?"
   b. Grace : "I am tired".

Grace's response is not immediately relevant with respect to Alice's question. So Grace's answer seems to be violating the maxim of Relation: what she literally says is not immediately related to the content of Alice's utterance, and cannot therefore count as an appropriate answer to Alice's question. Still, suppose that Alice has no reason to think that Grace is not cooperating to the conversation. In light of this datum, Alice should be able to run the following inference (which I adapt from Grice 1967/1989): "Assuming that Grace is being cooperative, she could not have uttered what she uttered unless she thought that she doesn't want to go to the movies; she knows that I can derive this supposition given what she said; she has done nothing to stop me from thinking that she thinks that she doesn't want to go to the movies; she wants me to think that she doesn't want to go to the movies; therefore, she is implicating that she doesn't want to go to the movies".

Another way an implicature might be generated is through explicit violation or "flouting" of one of the maxims. Such a blatant violation should trigger in the hearer a chain of reasoning similar to that reproduced above, of course provided that the hearer has no reason to doubt that the speaker is still being cooperative. A notorious example of maxim-flouting is irony: I could comment on my friend X's having betrayed me by saying "X is a fine friend". By using this sentence, I would be patently violating the maxim of Quality, for I would be uttering something strictly speaking false. Still, if my hearer regards me as cooperative, she could engage in the following train of thought: "I have no reason to think that she is not being cooperative; the only reason for her to utter what she uttered is that she thinks that X is an awful friend; she knows I can derive this supposition based on what she said; she has done nothing to stop me from thinking that she thinks that X is an awful friend; she wants me to think that X is an awful friend; therefore, she is implicating that X is an awful friend".

In general, when speaking about the process of implicature-calculation, it is taken as understood that the hearer has to go through what is being said (in virtue of its linguistic meaning) by a sentence first, in order to subsequently understand what is being implicated by the act of uttering that sentence. This becomes clearer if we consider the step-by-step reasoning that, according to Grice, leads to the derivation of an implicature. Suppose Grace asks Alice whether Charles is good-looking. As a response, Alice utters "Charles is a very smart person". The answer clearly is off-topic, in that it doesn't provide information about the way Charles looks, in contrast with what was required by Grace's question. In Grice's terms, the maxim of relation is violated, in that Alice's answer as to the intelligence of Charles is totally unrelated with the question of Charles' look. How could then the hearer of Alice's answer make sense of her utterance? Grice's answer is that the hearer may go through a reasoning such as those exemplified above, whose first premise is:

[1] Alice said that Charles is a very smart person;

and whose conclusion is:

[n] She is implicating that Charles is not good looking.

As one can note, the *first premise* of the reasoning is one in which the notion of *saying* occurs. Even though at this stage the notion could be totally theoretically unladen, it is worthwhile to note that, in order to derive what she meant to communicate, Grace cannot but go through what Alice conveyed by using the means that she and Grace are most likely (at least in their circumstances) to share: the English language. Implicature-calculation is therefore a reasoning in which one of the premises (usually, the first) has the form "X said that that *p*", where the notion of saying seems tightly linked with that of *linguistic meaning*. So, if the notion of implicature is ever plausible, then it requires that something is said in the first place, where the notion of "saying" is intimately related with the meaning expressions have in virtue of linguistic conventions.

2.2 Undermining Grice's Distinction
Some authors have questioned Grice's distinction between what is said and what is implicated, on the account that the notion of what is said, which is essential in making sense of the concept of implicature, is hardly ever exemplified: There is hardly ever anything around corresponding to what is *said* in Grice's sense.

In Travis' (1985/2008) interpretation of Grice, the notion of what is said is composed by two elements, which we might call a "conventional" element and a "truth-conditional" element. What does this mean? When an utterance of a sentence is performed, what is said is, in Travis' reading of Grice: (i) the result of the application of linguistic conventions, and (ii) what is "strictly speaking true", i.e. what in the utterance would be true or false, regardless of the implications or presuppositions raised by the utterance itself. Travis (1985/2008) illustrates (ii) with the help of an example made by Grice himself: what is said by an utterance of "He is an Englishman; he is therefore, brave" is that a certain individual is an Englishman and is brave. An utterance of such sentence would normally raise the implication that the man's being brave is a *consequence* of his being an Englishman. However, Grice argues, even if the man's being brave didn't follow from his being English, the sentence would still be *strictly speaking* true: "Grice takes it that the relevant person's being an Englishman *and* being brave is enough to make what was said, strictly speaking, true […]." (Travis 1985/2008: 22) [7].

At first sight, the conventional and the truth-conditional component seem to live together quite well. However, Travis points out, difficulties are behind the corner. Consider an utterance of the sentence "The leaves are green". If this utterance were to implicate anything, it would have to *say* something in the first place. In order to determine what the sentence says, we should be able to identify what makes the sentence strictly speaking true. However, determining such a thing is not straightforward at all. To illustrate the case, consider the following situation (inspired by Travis (1997)):

> Case A: Pia has just bought a ficus to be placed in the living room. The leaves of the plant have a characteristic dark and shiny green shade. Contemplating her new purchase, Pia utters "The leaves are green". The *ficus' leaves'* being green makes the sentence Pia uses true.

> Case B: Pia owns a Japanese maple, whose leaves are characteristically red. One day, Pia decides that she would like her Japanese maple's leaves better if they were green; therefore, she buys green dye and colours them green. Once she has finished, and

---

7   Travis's way of reconstructing Grice depicts him as very much in the footsteps of Frege, who expresses a very similar point of view about "and" and "but". He writes: "The word 'but' differs from 'and' in that with it one intimates that what follows is in contrast with what would be expected from what preceded it. Such suggestions in speech make no difference to the thought. A sentence can be transformed by changing the verb from active to passive and making the object the subject at the same time. In the same way the dative may be changed into the nominative while "give" is replaced by "I receive". Naturally such transformations are not indifferent in every respect; but they do not touch the thought, *they do not touch what is true or false*" (Frege 1956 p. 295-6, my italics).

while contemplating her work, Pia utters "The leaves are green". The *painted leaves'* being green make the sentence Pia uses true.

It seems that what is said by the sentence "The leaves are green" is different in Case A and in Case B. Why? Because in each case, what makes the sentence true is a different way for the leaves of being green. The property of being green in Case A, such that it makes *that* particular utterance of "The leaves are green" true, is different from the property of being green in Case B. So it seems that in the two cases, the same sentence "The leaves are green" says two different things. So, determining "what makes the sentence strictly speaking true" is of no help for determining what the sentence says. Since the sentence is made true by different states of affairs in different circumstances, what is said with that sentence changes across those circumstances. If there are many things that "can be said" with a sentence in different circumstances, there's nothing such as what that sentence "says". As Travis puts it: "there are at least several distinct things to be said […] in one or another speaking of our sample words […]. If so, then it is incorrect to speak of "what is said" […] in speaking those words" (Travis 1985/2008: 26).

However, one might object, taking "what makes a sentence true" as a guide for determining "what a sentence says" is misguided. After all, sentences have meanings! So, what is said by the sentence is presumably determined by the sentence's conventional *meaning*. For example, what "The leaves are green" says shall be determined by its conventional meaning, which yields the content *<that the leaves are green>*. But, again, here Travis notes a difficulty: what makes the simple semantic content *<that the leaves are green>* true or false? The meaning of the sentence does nothing to make what the sentence says capable of being true or false in any circumstance, in that it fails to exhaustively indicate what should be the case in order for an utterance of that sentence to be evaluable. For example, one may ask in what way the leaves would have to be green in order to count as such: would they have to be green on the outer surface? Or beneath the surface instead? And so on. As a consequence of these considerations, Travis concludes that sentence meaning is not what is said by the sentence, for it doesn't say anything such that could be true or false. "When what is said may be true or false, then, what is said in speaking given words cannot be revealed by their meanings […] on that speaking alone" (Travis 1985/2008: 33).

In sum, the reason why the Gricean distinction between saying and implicating doesn't stand is that the very Gricean notion of what is said is a non-starter: on the one hand, we cannot determine "what is said" on the basis of "what makes the sentence strictly speaking true", for "what makes a sentence strictly speaking true" changes from context to context, thus making it that there is no one single thing that the sentence, all by itself, says; on the other hand, we cannot determine "what is said" on the basis of conventional meaning, because meaning turns out as seriously under-specific.

2.3. Beyond the Gricean Distinction: Pragmatic "What is Said", Explicature and Impliciture

If Travis is right, and there's no way to determine "what is said" by a sentence, then the Gricean picture happens to be under threat. If a sentence doesn't say anything in virtue of its conventional semantic properties, then what is said will have to be determined by something other than convention. In the view of Travis and others (see below), *extra-linguistic*, pragmatic factors intervene in determining what is said. If what is said is determined by elements that do not pertain to the semantic conventions of language, then Grice's initial distinction between what is said and what is implicated gets blurred.

Those theorists that are happy to undermine the Gricean sharp distinction between "what is said" and "what is implicated" are all keen on the idea of there being a *third layer* between the sentence's conventional meaning and what is implicated by utterances of that sentence. This third layer still provides the truth-conditions of the uttered sentence, even though it is somewhat richer than the simple sentence's meaning: it is, as it were, a level for "pragmatically enhanced" sentence-meaning. Recanati (1989, 2001) calls this level "what is said" and characterises the notion in opposition to

both sentence meaning, which is too under-determined to deliver a propositional (i.e. evaluable) content, and to implicature, which is an altogether separate proposition. Thus he writes:

> The distinguishing characteristic of sentence meaning (the linguistic meaning of the sentence-type) is that it is conventional and context-independent. Moreover, in general at least, it falls short of constituting a complete proposition, i.e. something truth-evaluable. In contrast, both 'what is said' and 'what is implicated' are context-dependent and propositional. The difference between 'what is said' and 'what is implicated' is that the former is constrained by sentence meaning in a way in which the implicatures are not. What is said results from fleshing out the meaning of the sentence (which is like a semantic 'skeleton') so as to make it propositional. (Recanati 2001: 76)

Similarly Carston (1988, 2002, 2004), who largely shares her view with Sperber and Wilson (1986/1995), writes the following about the distinction between the sentence's linguistic meaning, implicature and a third, intermediate level of content, which she calls "explicature":

> An *explicature* is a propositional form communicated by an utterance and is pragmatically constructed on the basis of the propositional schema or template (logical form) that the utterance encodes; its content, therefore, is an amalgam of linguistically decoded material and pragmatically inferred material. An *implicature* is any other propositional form communicated by an utterance; its content consists of wholly pragmatically inferred matter. (Carston 2004, 824)

As for Bach (1994a, b), he names the third level "impliciture". The notion he employs, though still related to the ones characterised in the quoted passages above, is slightly different from that put forth by Recanati and Carston. While Recanati and Carston have it that the sentence's meaning remains implicit, while what is explicitly conveyed, through a combination of linguistic and pragmatic elements, is the augmented "what is said" or "explicature", Bach wishes to emphasise the *implicit* character of this third level of significance. This stress on implicitness depends on the way Bach conceives the contents of sentences. In Bach's view, the content expressed by the utterance of a sentence is akin to a *structured proposition,* i.e. a sentence-like entity, consisting in objects and properties bound by appropriate relations. Structured propositions typically maintain or at least faithfully reflect the structure of the sentences that express them. If propositions are thus constrained to the structure of sentences, then it follows that, at the explicit level, no elements that are extraneous to the sentence's structure can be included in the proposition expressed by the utterance. Rather, these extraneous elements belong to the *implicit* part of the utterance's content, i.e. that part of the content which is left verbally unexpressed. This entails that an utterance of a sentence *s*, which explicitly expresses the proposition *p*, may *implicitly* express a different proposition *q*. This proposition is what Bach calls "impliciture".

Apart from these technical details, an impliciture is, just as Recanati's "what is said" and Carston's explicature, the result of a pragmatic augmentation of sentence meaning, which for Bach is often semantically under-determined. Thus he writes:

> Implicitures go beyond what is said [NB: here by "what is said", Bach intends "sentence meaning"], but unlike implicatures, which are additional propositions external to what is said, implicitures are built out of what is said. Even when there is no figurative use of words or phrases, as in metaphor, in impliciture [...] what the sentence means does not fully determine, even after ambiguities are resolved and indexical references are fixed, what the speaker means. (Bach 1994b: 273)

To sum up, the demotion of the Gricean paradigm has given impulse to an alternative view, which has turned the Gricean dichotomy into a three-layers partition: at the first level, there is meaning, which is determined by linguistic conventions, but often fails to deliver the utterance's truth-conditions; at the second level, there is "pragmatically enhanced" meaning, i.e. semantic content appropriately enriched, or expanded in response to pragmatic demands; at the third level, the Gricean notion of "what is implicated" is preserved, as a level of significance which radically departs from the utterance's truth-conditions—even the pragmatically augmented ones.

2.4 A Challenge to Compositionality

To accept, as Travis does, that the meaning of some sentences fails to specify the truth-conditions of utterances, seems to lead to a view in which the content of these sentences is determined by a combination of semantic *and* pragmatic processes. A consequence of this is that the boundary itself between the "semantic domain" and the "pragmatic domain" gets blurred. Theorists have looked at this conclusion as *a challenge to the idea of the compositionality of utterance content* (see Stanley, 2000, 2002, 2005b, Stanley & Szabó 2000, King & Stanley 2005). In this section, I will introduce some views which contrast the idea that a sentence's content in context is compositional and explain why their proposal is problematic.

Let us accept the following formulation of the Principle of Compositionality:

[The Principle of Compositionality] The content of an utterance *u* of a sentence *s* is determined (*modulo* ambiguity resolution and contextual reference-assignment) by the semantic content of the components of *s* and their syntactic combination.

In a standard formal semantic framework, the input of the semantic interpretation of a sentence *s* is an "interpreted logical form", i.e. a complex entity which contains the meanings of the sentence's constituents (such as they are found in the lexicon) and the sentence's syntactic form, while the output corresponds to the sentence's truth-conditions. Suppose one has a fragment of the English language L which contains only the proper name "Spot" and the predicate "is happy". "Spot" denotes the individual *o* and "is happy" denotes a function form individuals to truth-values, which maps onto truth the individuals that are happy and onto falsity the individuals that are not happy. The semantic interpretation of a sentence like "Spot is happy" will proceed from the compositional semantic rule to the effect that, whenever a sentence is composed by a name *"α"* and a predicate *"φ"* in the form *"φ(α)"*, then the sentence will be true if and only if the reference of the name, *α*, is in the extension of the predicate *φ*. In the case at hand, the sentence "Spot is happy" will be true iff the individual *o* which is referred to by the name "Spot" is in the extension of the predicate "is happy" (i.e. iff the function denoted by "happy" maps *o* onto truth). This way of determining the interpretation of a string of symbols is characteristic of a *compositional semantics*, i.e. a semantics according to which the meaning of a complex expression is determined by the meanings (in context) of its components and their syntactic composition.

Now, among the theorists who accept semantic under-determinacy, some may feel attracted to the view that not only semantic processes, but also pragmatic processes contribute to the determination of the content of utterances of these expressions, where this implies that compositional rules will not be the *sole* determinants of the utterance's content. In other words, those who are happy to admit an intrusion of pragmatic processes into semantic interpretation may be willing to commit to the idea that utterance content is not determined compositionally.

As we argued in the **Introduction** and illustrated in **section 1**, colour predicates such as "green" or "red" might be considered as giving rise to semantic under-determinacy. Consider for example the phrase "Red pen". If we abide by the principle of compositionality, we have to say that the meaning of "Red pen" is determined by the meaning of "red" and the meaning of "pen" in the way illustrated

above; here the way the predicate "red" contributes to the overall phrasal meaning may be such that the phrase results as expressing something like *<pen with red ink>*. Now the same phrase in a different context may express a different content, perhaps something like *<pen made with red plastic>*. If "red" means the same in both phrases, how come that the overall interpretation is different? According to Lahav (1989), this variation in the truth-conditional contribution of "red" suggests that its semantics (and the semantics of similar adjectives) is non-compositional: that is, in the interpretation of a complex expression that contains an occurrence of "red", the word's meaning gives only a partial contribution. There must be some additional, intrinsically non-compositional (probably pragmatic) mechanism which operates in combination with the rules of compositionality in delivering the content of each of the different complex expressions in which the word "red" occurs. Similar points about the semantics of quantified phrases like "Every students", are advanced in Pelletier (1994, 2003).

If theorists such as Lahav and Pelletier are right, then at least when it comes to the semantic workings of certain groups of expressions, such as colour adjectives and quantifiers, it is legitimate to say that the content and truth-conditions of sentences that contain these terms are not determined compositionally; non-compositional and plausibly pragmatic mechanisms "intrude" in the determination process. If that is so, a truth-conditional semantics project such as that set out by Frege (1893/1903), or in Wittgenstein's *Tractatus* or in Tarski (1936), and applied to natural language by philosophers such as Davidson (1967a), Lewis (1970), Montague (1974), is mistaken, in that it fails to capture the real semantic functioning of some expressions. The central assumption shared by all these theorists is that meaning determines truth-conditions, where the word "determines" means that a sentence's truth-conditions are determined *uniquely* by the meaning (in context) of the sentence's components and their syntactic combination. If there is semantic under-determinacy, however, meaning doesn't strictly speaking determine truth-conditions, and something other than meaning—e.g. pragmatic mechanisms—must enter the scene. The truth-conditional semantics projects gets therefore undermined.

But would it be such a big deal if truth-conditional, compositional semantics were mistaken? There's no need to be dogmatic here in order to see that a compositional semantics has virtues that should be regarded as valuable in their own right. I take it that, primarily, a compositional semantics helps make sense of that part of our communicative behaviour that falls under the heading of *linguistic* behaviour. It sets language apart from the rest of our communicative behaviour, and captures what makes language special as opposed to gazes, gestures, vocal signals, etc. Thus, it is important to maintain compositionality to the extent that it is important not to loose sight of the peculiarity of language as opposed to other devices of communication human beings have at their disposal.

A further motivation for not abandoning a compositional account (advocated by Stanley, 2002), is that, if we accept the contribution of pragmatic processes that are not linguistically controlled, and hence which do not follow compositional rules, then we have no means to constrain the effects of these processes. This point rests on the assumption that pragmatic effects can be "constrained" only if they are tied to syntactic structures, which of course respond to compositionality. If syntactic constraints are given up, then virtually any pragmatic mechanism can operate on content—even mechanisms that hearers cannot actually manage. For instance, it seems impossible that a hearer could pragmatically supplement the content of an utterance of "Everyone likes Sally" so as to obtain *<everyone likes Sally and his mother>*. However, if we abandon a syntactic (compositional) approach, it seems that we have no means to explain why this is not feasible. In conclusion, a compositional semantic model should be preserved, on pains of losing a way of constraining the effect of pragmatic processes.

To sum up, acknowledging semantic under-determinacy may lead the theorist to regard utterance content as also pragmatically determined. This has been considered as an attack to compositional semantics. Compositionality is valuable in itself and also for the constraints it poses on "wild"

pragmatic effects. For these reasons, some theorists have argued, we'd better not forsake compositionality.

3. The Demarcation Problem
In the preceding sections, I introduced the notion of semantic under-determinacy, initially by simply fostering some pre-theoretic intuitions and subsequently by framing the problem historically. I have gone through Grice's saying-implicating distinction and Travis' rejection of the distinction, on the account that there is no such thing as "what a sentence says". I subsequently introduced theories which supplement the Gricean saying-implicating dichotomy with an intermediate level of content, which is called by Carston "explicature", by Recanati "what is said" and by Bach "impliciture". This intermediate layer of content consists of the sentence's compositionally determined semantic content (in context) *plus* additional material which is supplied by pragmatic processes. Finally, I have illustrated how theorists who welcome this additional level of content may be charged of rejecting the principle of compositionality—where this means rejecting the idea that the content of a sentence (in context) is determined *exclusively* on the basis of the meaning of its components and their syntactic arrangement.

Having introduced the notion of semantic under-determinacy and its implications, in this section I would like to concentrate on a question which has been tackled only marginally by those theorists who have been engaged with this idea. The problem concerns drawing a sharp demarcation between cases of semantic under-determinacy and cases of semantic under-specificity which are clearly due to ambiguity, vagueness, indexicality, ellipsis and so on. I will call this *The Demarcation Problem*, which could be thus formulated:

[The Demarcation Problem] Is there a feature of semantically under-determined expressions which sets them apart from other cases of semantic under-specificity, such as ambiguity, vagueness, indexicality, ellipsis, etc.?

I believe that providing a satisfying answer to the Demarcation Problem is a key step for any theorist who wishes to understand the nature and extent of semantic under-determinacy in natural language. Such an answer is particularly called for also in view of the fact that (as I will illustrate in **section 5**) theorists have so far only succeeded in saying what semantic under-determinacy is *not*—not ambiguity, not vagueness, not indexicality, not ellipsis (see Searle 1980, Travis 1975, 1997, Bezuidenhout 2002). This however, doesn't yet license any claims as to what semantic under-determinacy *is*. In this chapter, my aim is that of saying something more than simply what semantic under-determinacy is not, even though I won't engage in the attempt of saying any final word about what semantic under-determinacy is: that is, I will try to do best than merely negatively characterise semantic under-determinacy, even though I won't aim at providing any necessary and sufficient conditions for its identification. Providing an answer to the Demarcation Problem is a compromise between the unsatisfying result of characterising semantic under-determinacy only negatively (by saying what it is not) and the—probably too ambitious—attempt to *define* semantic under-determinacy. In order to answer the Demarcation Problem, it will be sufficient to point at at least *one feature* which semantically under-determined expressions have which is not shared by ambiguous, vague, indexical or elliptical expressions. If there's at least one feature which semantically under-determined expressions have and which these other expressions do not have, then the Demarcation Problem will have been answered to. Answering the Demarcation Problem will not coincide with providing a definition of semantic under-determinacy insofar as the feature at issue needs not be necessary and sufficient—it may just be necessary.

The rest of this chapter will be devoted to providing an account which can fulfil the need for an answer to the Demarcation Problem. The tasks I will undertake in the next sections will therefore be the following: first, I will reconstruct the main argument that has been provided so far in order to

draw attention on semantic under-determinacy (**section 4**); then I will move to the arguments so far advanced in order to set semantic under-determinacy apart from other, apparently familiar phenomena: as it will emerge, all these arguments just succeed in telling us what semantic under-determinacy is *not* (**section 5**). In **sections 6-9**, I will try to pin-down the source of semantic under-determinacy by introducing the notion of under-articulation and, subsequently, I will set forth an answer to the Demarcation Problem which employs this very notion. In **sections 10**, I will confront an important objection to the proposal.

4. The "Context-Shifting" Argument

The considerations set forth in the **Introduction** were aimed at providing a provisional way of restricting the scope of semantic under-determinacy to those under-specific expressions which, however, did not immediately strike us as examples of ambiguity, vagueness or context-dependency. Indeed, delimiting the scope of application of the word "semantic under-determinacy" is not an easy task. Triggering intuitions as to the semantic defectiveness of sentences like "It's raining", "Jill is ready", "The leaves are green" etc. (see **sections 1** and **2.2**) is not equivalent to providing an argument for their semantic under-determinacy. Intuitions of defectiveness surely point in the direction of *some* under-specification, but that is still too unqualified: the under-specification could be explained as a form of ambiguity, vagueness, indexicality and so on. Proponents of semantic under-determinacy therefore need an argument in order to establish that the under-specification they are pointing at is not an instance of any other familiar phenomenon. The typical argument to this end is what has been called "The Context-Shifting Argument" (CSA)[8] (see Travis (1975, 1985/2008, 1996, 1997) and Bezuidenhout (2002)). The argument could be thus reconstructed:

CSA
[1] Suppose a sentence s is uttered within a context c_1, where it turns out true;
[2] Suppose the same sentence s is uttered in a second context c_2, where it turns out false;
[3] There is no *prima facie* reason to think that s contains any indexical, ambiguous, vague components, nor that it involves sentential ellipsis;
[4] Between the first and the second utterance, no relevant change in the way the world is has occurred;
[5] The meaning of s semantically under-determines its utterances' truth-conditions.

Let us illustrate how [5] is obtained from [1]-[4]. The argument has the structure of an inference to the best explanation. The question might be put this way: in light of the conjunction of [1]-[4], how is the truth-value shift of s to be explained? For these authors, the explanation lies in the linguistic content of s's being in defective in some peculiar way. For, as they reason, if the truth-value of s varies across contexts, even though how the world is doesn't change, and there is no reason to suppose that s exhibits any form of indexicality, ambiguity, vagueness, ellipsis etc., this means that the linguistic content of s is not determined enough to give the truth-conditions of its utterances in c_1 and c_2. To illustrate how the argument works, here is Travis's famous green leaves example:

> Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, "That's better. The leaves are green now". She speaks truth. A botanist friend then phones, seeking green leaves for a study in green-leaves chemistry. "The leaves (on my tree) are green" Pia says. "You can have those". But now Pia speaks falsehood. (Travis 1997: 89)

In this example, the same sentence can acquire different truth values, even though no relevant changes in words and worldly conditions occur (the maple leaves keep being painted green both on

---

8   This term is due to Cappelen and Lepore (2005).

the first and on the second occasion). So, when uttering "The leaves are green" the first time, Pia utters something true, whereas she utters something false while she speaks the same words with her botanist friend. Now there's no *prima facie* reason to think that that particular utterance of the predicate "being green" exhibits ambiguity, indexicality, vagueness, nor that the sentence "The leaves are green" is elliptical. In view of these considerations, that one single sentence can shift its truth-value across contexts in this way is explained by the fact that:

> What [...] words mean leaves it open for them to be used (in suitable circumstances) to say any of various things, each true under and on, different conditions. There is no one set of conditions under which those English words […] would be, or say what is, true. Nor even one condition which is the condition for them to be true. If differences in truth condition make for different propositions, then what those words mean makes no one proposition the one (modulo referents) they express. (Travis 1996: 454-455).

Similar remarks—only with slight changes in terminology—can be found also in other authors. In the following quotes, both Carston and Recanati give voice to the same idea:

> The linguistic semantics of the utterance, that is, the meaning encoded in the linguistic expressions used […] under-determines the proposition expressed (what is said). (Carston 2002: 19-20).

> Semantic interpretation, characterized by its deductive character, does not deliver complete propositions. (Recanati 2004: 91).

Context-Shifting Arguments can be regarded as successful only if their premises are adequately warranted. With regard to this demand of warrant, the premise that stands most in need of justification is certainly [3], which states that there is no *prima facie* reason to think that a sentence like, for instance, "The leaves are green" contains any indexical, ambiguous, vague components, nor that it involves sentential ellipsis. This claim needs to be sustained by detailed, case-by-case arguments aimed at ruling out that semantic under-determinacy is identified with any other phenomenon of semantic defectiveness. I shall turn to a reconstruction of these arguments in the next section.

5. No Domestication
As we noted earlier, proponents of the semantic under-determinacy thesis make much of the assumption that there is no (*prima facie*) reason to think that what is going on in examples like Pia's story is a manifestation of some phenomenon like indexicality, ambiguity, vagueness, etc. In their view the under-determinacy shown by Pia-style cases is not reducible to any phenomenon of under-specification already familiar in the study of natural languages. In this section I shall survey the motivations that have been advanced to support this idea.
[1] No ambiguity: It has been argued that semantic under-determinacy cannot be a case of ambiguity. First, the semantic variability typical of semantic under-determinacy and of ambiguity is different in kind: words like "green" have indefinitely many senses depending on context, while ambiguous words have definitely many senses, independently of the context in which they are used (see Bezuidenhout 2002: 107-8). Secondly, I take it that the necessary conditions for lexical ambiguity are that (1) there be a single lexical entry to which more than one meaning is assigned by the dictionary; (2) the number of meanings, although it can increase or decrease as time goes by, has to be definite, for reasons connected to language acquisition and memory. Let us then consider a word like "green". One could surmise that the lexical entry "green" is associated with more than

one meaning *green₁, green₂, green₃ … greenₙ*, which gets selected according to the context of utterance. To this idea one can object that, since there are indefinitely many ways of being green, the word "green" will have to be ambiguous between indefinitely many meanings. This, however, goes against the necessary requirement (2) for ambiguity. Thirdly, take a typically ambiguous word like "bank". If Sam goes to the bank (in the sense of a financial institution) and Sal goes to the bank (in the sense of river side), then, once "bank" is disambiguated in context, it is false to say "Sam and Sal both went to the bank". This however, wouldn't happen with "green". Suppose Pia painted the leaves green, making them green on the outside; Lia, on the other hand, brings a green parrot. Even in light of these differences between the ways of being green of the leaves and the parrot, it would seem correct to state "The leaves and the parrot are both green" (see Searle 1980: 224).

[2] No Indexicality: Travis makes the following case, speaking from a Fregean point of view on indexicality: First of all, it doesn't seem part of what "green" means that it points at a parameter, just like "I" points at the speaker-parameter or the present tense points at the time of speaking. Secondly, if the speaker were to hear an utterance of "I am in Paris", but was in the dark about what the context of utterance was, she could still know what state of affairs the meaning of the sentence describes: namely, that the speaker is in Paris. This is because it is part of the meaning of "I" that, whatever its referent, it shall always be identical to the speaker. The same doesn't hold for "The leaves are green": upon hearing an utterance of "The leaves are green", while being in the dark about the context of utterance, the hearer cannot know *what state of affairs* is described by the utterance because "green" doesn't point to any fixed parameter and so, unless more information is provided, the sentence doesn't describe *any* state of affairs (see Travis 1997: 91-94).

A further argument against the equation of semantic under-determinacy with indexicality can be read in Searle (1980) (and also in Travis 1997, but somewhat between the lines): suppose that a term like "green" is associated with a free variable that takes different values in different contexts, thereby making the word express a different property on each occasion. So for example, when "green" takes as argument *on the outer surface*, it expresses the property of being green on the outer surface, while if it takes as argument *underneath the paint* it expresses the property of being green underneath the paint. This move, it is argued, would not make the truth-conditional contribution of "green" determined, even once the argument has taken a value in context. For suppose the value of the variable were fixed in context c_i so that "green" expressed the property *<being green on the outer surface>*. There would be *n* ways for taking an item to be green on the outer surface: for example, because a green light shines on it, or because, even though the item is not exclusively green on the outer surface, for present purposes it is convenient to call it "green" just in case it is green on a sufficiently great part of its surface, or because, even though the item is not entirely green on its outer surface, it looks entirely green when looked at from far away... and so on. So, fix the (putative) variable value of "green" in c_i, thereby resolving indexicality, and still it will be an open question which property the predicate expresses in c_i, thereby leaving the semantic under-determinacy of the predicate unaltered.

Other considerations against the indexicality proposal are voiced by Bezuidenhout: first of all, she says, the proposal sounds *ad-hoc*, in the sense that "it posits a hidden variable not for any well-motivated syntactic reason, but simply because it wishes to subsume [semantic under-determinacy] under the heading of something allegedly better understood." (2002: 113). Moreover, since semantic under-determinacy is highly pervasive, this entails that the theory "would have to posit syntactic representations positively bristling with hidden indexicals" (2002: 113), and the proliferation of hidden indexical elements is not desirable neither for the sake of theoretical economy and elegance, nor for its cognitive plausibility.

[3] No vagueness: semantic under-determinacy cannot be vagueness either. Even though the matter of what vagueness is is itself controversial, let us assume, in line with the standard conception, that a necessary feature of vagueness is the obtaining of borderline, or unclear cases (e.g. of baldness, tallness, fatness, greenness etc.). Now one should of course concede that an expression like "green"

does give rise to borderline cases—after all, there are many cases in which, by looking at an object, we cannot tell whether it is green or yellow, or green or blue. However the particular defectiveness that makes "green" semantically under-determined is arguably not connected with its being vague. To illustrate, consider an utterance of "The leaves are green" which one assesses as not determinately true or false (maybe because the leaves are of an intermediate shade between green and yellow); in order for one to be able to say that the utterance doesn't have a definite evaluation in the first place, one needs to *already* have figured in which respect the salient leaves are to be said green—on the outside, underneath the surface, etc. That is, in order to judge about the vagueness of a statement of "The leaves are green" one must previously resolve the semantic under-determinacy of the sentence's content as to the respect in which the leaves are green. (The same, of course, holds also for cases in which an utterance of "The leaves are green" has a definite truth-value. So, any judgement concerning truth value requires a prior resolution of semantic under-determinacy.) This seems to make a case for the idea that semantic under-determinacy is not reducible to vagueness, on the account that acknowledging vagueness implies, at least in some cases, resolving precisely semantic under-determinacy (for similar considerations, see Travis 1997: 91; Bezuidenhout 2002: 115).

[4] No ellipsis: It could be suggested that semantic under-determinacy is just sentential ellipsis. But this is implausible, too. For suppose that "The leaves are green" were elliptical for some longer sentence, namely "The leaves are green underneath the paint". How is this supposed to be found out by the listener of an utterance of the former sentence? Elliptical sentences usually fail to articulate linguistic material which was articulated in some previous utterance. An utterance of an elliptical sentence is typically *infelicitous* if the audience is not in a position to recover the elided material from previous speech. By contrast, it seems that "The leaves are green" could be felicitously uttered even if no previous sentence has been articulated before that (provided that the utterance has a sufficiently rich background of presuppositions, for example as to what part of the leaves is being discussed). This seems a good reason to rule out that semantic under-determinacy is just a case of sentential ellipsis.

I believe that the points raised at *[1]-[4]* are on the right track. Semantic under-determinacy should not be reduced to any already familiar phenomenon of semantic defectiveness like ambiguity, indexicality, vagueness, ellipsis and so on. In what follows, I will act on the assumption that a *no-domestication policy*, as I shall call it, is the way to go if one wants to make sense of the peculiarity of semantic under-determinacy as opposed to other phenomena of semantic defectiveness already present in language. Still, I believe that the no domestication policy isn't enough to provide an account of what makes semantic under-determinacy different from *all* the phenomena thus far surveyed. That is, adoption of the no-domestication policy doesn't yet license the theorist to say what distinguishes semantic under-determinacy from all the other cases of semantic defectiveness, for it doesn't point at any feature which semantically under-determined expressions have which all other expressions lack. What in **section 3** I called the *Demarcation Problem* still needs an answer. As it will appear in the following sections (especially **sections 8**-**11**), my attempt will be that of differentiating the phenomenon of semantic under-determinacy by identifying its *source*. I will attempt to capture what makes semantic under-determinacy peculiar and different from the phenomena listed above by exploring the issue of what more basic phenomenon might lie at the origin of it. Doing this will enable me to make sense of the difference between semantic under-determinacy and other phenomena like indexicality, ambiguity, vagueness and ellipsis.

6. What is the Source of Semantic Under-determinacy?
As we have seen, authors such as Searle, Travis, Bezuidenhout, Recanati, Sperber and Wilson and Carston all have subscribed to the claim that sentences like "The leaves are green", "Jill is ready", "It's raining" are semantically under-determined. Moreover, their view commits them to saying that semantic under-determinacy is irreducible to any other already familiar phenomenon of under-

specification such as ambiguity, indexicality, vagueness or ellipsis. Many arguments are available to the theorist who wishes to negatively characterise semantic under-determinacy—i.e. to say what semantic under-determinacy is not. Having established this much, there is one residual question that is left open, corresponding to the *Demarcation Problem*: is there a feature which distinguishes semantic under-determinacy from all these other familiar phemomena? In what follows, I will attempt to deal with the *Demarcation Problem* by exploring the issue of what lies "at the source" of semantic under-determinacy. The first step in this direction will therefore imply asking ourselves a question like the following: in virtue of what are semantically under-determined expressions such? I believe that the following are the most spontaneous answers to the previous questions:

Answer (1)
Sentence meaning under-determines utterance content (and truth-conditions) because *meaning itself is indeterminate*;

Answer (2)
Sentence meaning under-determines utterance content (and truth-conditions), but meaning itself is determined: it just fails to determine the *content-in-context* of utterances.

I will deal with each of these answers in the following two sub-sections.

6.1 Indeterminacy of Meaning

Let us begin by considering **Answer 1**. Stated as it is and with no further qualifications, it is a very strong answer. Saying that meaning is indeterminate means endorsing a view at the metaphysical level: that is, it means maintaining that nothing in the world can settle the question of what the meaning of a certain expression is. Throughout the history of analytic philosophy, different arguments have been invoked to support this thesis. The most famous is probably the one mounted by Quine in support of his thesis about the indeterminacy of translation. What Quine takes to have shown with his famous example from radical translation (see Quine, 1960, 1968) is that for any sentence, no matter if in our own mother tongue or in the most remote foreign language, it is possible to provide an indefinite number of translations that fit the relevant objective facts about language use and yet are incompatible with each other. The upshot of these remarks is that there is no fact of the matter as to which is the correct translation of the sentence at issue; this ultimately entails that there is no fact of the matter as to what the sentence at issue means, nor as to what its truth-conditions are: both intension and extension are indeterminate, or inscrutable. Another influential argument that brings to the conclusion that meaning is indeterminate can be found in Kripke's (1982) interpretation of Wittgenstein's rule-following argument in the *Philosophical Investigations*. The argument (as reconstructed by Kripke) goes like this: although one may be confident that, in performing a calculation like 7+5=12, one is using the "+" sign as denoting the "plus" function, in one's past one might as well have used "+" as meaning a completely different function—say, the "quus" function. How does one know one has never used "+" as denoting the "quus" function? There is no past or present fact one can point at that could help settling the question—neither pertaining to one's behaviour or pertaining to one's internal mental states. In other words, one's past (and present) actions and thoughts are compatible with having used "+" as meaning plus as well as quus. If this is the case (and if any speaker can be subject to this kind of challenge) then it follows that the meaning itself of "+", and that of any other linguistic expression,

is indeterminate[9], that is, no fact in the world can establish that an expression *e* has a certain meaning.

Unless some qualifications are provided, the semantic indeterminacy thesis, such as it is, is compatible with the idea that *meaning can be shaped by the intentions, beliefs, and communicative purposes of the speakers*. This view is suggested by the following passage by Brandom[10]:

> A word – 'dog', 'stupid', 'Republican' – has a different significance in my mouth than it does in yours, because and insofar as what follows from its being applicable, its consequences of application, differ for me, in virtue of my different collateral beliefs […]. [...D]ifferent interlocutors […] use different concepts, attach different meanings to their words. (Brandom 1994: 587-8)

The link between the thesis of meaning indeterminacy and the idea that meaning depends on the speakers' beliefs and intentions is easy to see: to the extent that there is no fixed semantic property that linguistic items have, there is room left for their users to confer a meaning to linguistic expressions, according to the conditions of application that seem most appropriate, depending on the collateral beliefs and intentions of the users themselves.

This position is subject to the following difficulty, which makes it ultimately unappealing as a theoretical option for the theorist engaged in a study of how communication works. It seems that whether a word, or any linguistic item, is used correctly or incorrectly depends on whether the user employs that linguistic item *for what it means*. In other words, the correctness or incorrectness of linguistic use seems to be based on the fact that terms have independent semantic properties. If **Answer 1** is true, and if no other qualifications concerning the nature of meaning indeterminacy are provided, then it is easy to fall prey of the thought that what decides the correctness or incorrectness of linguistic use is principles that are not immediately semantic, but are pragmatic in nature and pertain to convenience, informativeness, perspicuity with respect to the beliefs, intentions and communicative aims of the speaker. These principles do not respond directly to the semantic properties of the words; they rather "shape" these properties according to the occasion of use and according to what the speakers believe and are aiming at. This implies that, if it is correct from this purely "pragmatic" point of view that a speaker employs the word *w* as if it meant *p*, and if the use doesn't generate incomprehension between the speaker and his audience, then the use is correct. If this possibility is admitted, however, then the problem arises of how to distinguish between what we ordinarily consider as *strictly speaking correct* uses (i.e. uses that follow the "semantics" of the words) and uses that are not strictly speaking correct, but are merely *admissible* given the

---

9   Finally, Putnam could be regarded as another supporter of the meaning indeterminacy view, because of his endorsement of a position called "Internal Realism" (see Putnam, 1978): roughly, Internal Realism has it that the claim that there is a "correspondence" or "reference" relation between the words of a language and the world cannot be an absolute claim, but it can only be a theory-relative claim. Putnam defends this idea with a twofold argumentative strategy: on the one hand, he emphasises the unintelligibility of the hard-core realist claim to the effect that a theory with various epistemic virtues (completeness, consistency, correctness in prediction, simplicity, plausibility) might still be false. On the other hand, he shows that a weaker claim is true: that a theory with such epistemic virtues might be just one among the various theories with which we capture the world. This literally "dissolves" the issue of a theory's absolute truth/falsity, because it makes the semantic relations like "satisfaction", "reference" and "truth" ultimately theory-relative (i.e. epistemic).

10 But see also Bilgrami (1992). As he writes "it is unlikely that any two people will have the same concept of anything, since it is unlikely that they have the same beliefs associated with the term which expresses that concept. This is the level on which theories of meaning do their work: they specify the contexts and the term-meanings of an agent along the lines of this [...] method. [...] [T]he concepts are very fine-grained and they are hardly ever shared by people." (Bilgrami 1992: 11). Davidson 1986/2005 draws pessimistic conclusions as to relation between communication and convention: "We must give up the idea of a clearly defined shared structure which language-users acquire and then apply to cases. And we should try again to say how convention in any important sense is involved in language; or, as I think, we should give up the attempt to illuminate how we communicate by appeal to conventions." (Davidson 1986/2005: 265)

circumstances of the conversation: for example malapropisms, or non-literal uses. It may be that, deep down, this distinction doesn't make sense at all; however, distinguishing between uses that are semantically kosher and uses that are not, but are still admissible in another, more pragmatic and for-the-sake-of-conversation sense, would be at least a way of remaining faithful to common sense[11].

This distinction is particularly important for our purposes, because it helps one to make an at least intuitive case for the claim that semantic under-determinacy doesn't depend on meaning indeterminacy. Intuitively, when a speaker uses a sentence like "The leaves are green", or "It's raining", even though the sentence fails to describe a state of affairs in the world such that it would make that utterance true, nevertheless the intuition is maintained that the speaker is using those words *for what they mean*, i.e. the intuition is maintained that the speaker is conforming to the semantic conventions that govern the use of words such as "to rain" and "green". If the use is judged semantically correct, then meaning is (sufficiently) determined. But if meaning is (sufficiently) determined, then it is not meaning indeterminacy that causes semantic under-determinacy. Ultimately, then, at least at an intuitive level, it seems that the under-determinacy of these sentences does not depend on the instability of meaning itself, i.e. on meaning indeterminacy.

In sum, I believe one should not give too much credit to **Answer 1**, for two reasons: firstly, meaning indeterminacy with no further qualification easily leads to the view that meaning is determined by the beliefs and intentions of speakers, thus obscuring the difference between semantic correctness and pragmatic correctness. Secondly, usages of semantically under-determined expressions are generally judged semantically correct, i.e. they are usages of the words for what those words mean. If there can be a judgment of semantic correctness, then meaning is sufficiently determined, and meaning indeterminacy is not what determines semantic under-determinacy.

6.2. A Failure to Determine Content-in-Context

Let us now turn to **Answer 2**, which appears to be far less radical. Indeed, it admits that meaning is not indeterminate *in itself*. What suffers from indeterminacy is the *content-in-context* of a certain sentence or expression, something which is definitely compatible with the expression having a stable meaning. So, for example, according to this view, "green" does have a determined, standing meaning. However, for any utterance of "The leaves are green", the meaning of the sentence simply determines the content *<that the leaves are green>*. This, for the proponent of semantic under-determinacy, isn't enough for characterising which particular state of affairs makes the utterance true or false; hence, the meaning of the sentence (though stable) doesn't determine the utterance's truth-conditions.

This way of characterising semantic under-determinacy, though it appears to be much more plausible than the one suggested by **Answer 1**, makes those expressions that, by the lights of

---

[11] Incidentally, tracing a distinction between uses that are "semantically correct" and uses that are merely "pragmatically acceptable" could be regarded also as a way of opposing a holistic approach to meaning. This is because—as far as I can see—the idea that meaning depends on the speaker's beliefs and intentions, coupled with the thesis of meaning holism, gives rise to the idea of meaning indeterminacy. Very sketchily, meaning holism holds that in order to know the meaning of a word/sentence, one has to know language as a whole, in the sense of knowing the meaning of every other word/sentence in the language. If different speakers of the same language L have different beliefs concerning the meaning of words/sentences, this will entail that a single word/sentence in L can have different meanings according to different speakers (who determine different "translation manuals"), the consequence being that there is no objective fact of the matter as to what the meaning of that word/sentence is. Among the proponents of meaning holism are Quine (1953), Brandom (1994), Davidson (1967a), Block (1995), Field (1977). Against holism, Michael Dumment proposes a position called *molecularism* (see Dummett 1991, chapter 10). The molecularist position has it that in order to know the meaning of a word/sentence of L, one has to know the meaning of some other words/sentences of a relevant fragment of L, but needs not know the meaning of all the words/sentences of L. The distinction between "semantically correct" uses and uses that are merely "pragmatically acceptable" may interestingly relate with a molecularist insight on meaning. Thanks to Annalisa Coliva for pointing this out to me.

Answer 2, are semantically under-determined, *very similar to indexical expressions*. Take for example a word like "I": It has a determined, context-independent meaning (which Kaplan called "character"). This character is merely a function that maps the context of utterance onto the speaker in that context, and therefore cannot be said to contribute the truth-conditions of *utterances* containing occurrences of "I". So, one could say that the meaning of, for instance, "I wear glasses" doesn't determine the truth-conditions of its utterances, because, unless a context is provided, the meaning of the sentence yields *too poor* a content in each context. This characterisation of the way "I" is semantically under-determined parallels the previous characterisation about "green": But is this parallel accurate? As we have seen, there are arguments available that show that indexicality and semantic under-determinacy are associated with different semantic results. Searle (1980) and Travis (1997), for example, both offer reasons in support of the distinction between semantic under-determinacy and indexicality, as illustrated in **section 4**.

Capitalising in part from Travis' and Searle's results, in this section I would like to suggest an argument in the form of a *reductio*, in order to defend the distinction between the semantic under-determinacy and the indexicality of certain terms. Since we already know what indexicality is, it will be convenient to assume that a certain, allegedly semantically under-determined word, say "green", is in fact indexical. This will make life easier to us, since we already know what the implications of being an indexical are and, on this basis, we will be able to highlight that these implications do not obtain in the case at interest. The needed premises of the *reductio* are: (1) the premise to the effect that some designated expression, say "green", is indexical, in that for example it features an extra covert argument in its logical form (the predicate's structure being *green(x, y)*); (2) a premise concerning a general feature of indexicality—that we will call the Indexicality Assumption:

[The Indexicality Assumption] Indexicals (are designed to) work in such a way that, once the reference of an indexical term i is fixed in a context k, then i expresses a determinate content, in every circumstance of evaluation.

The Indexicality Assumption may be viewed as stemming from the idea that indexicals are devices for contextual reference-assignment that are designed for being "determinately saturated". That is, indexicals are designed in such a way that, once an indexical expression i is assigned a referent according to the value taken by a certain parameter in a certain context k, the under-specification of the expression is *ipso facto* resolved in all possible circumstances of evaluation, in the sense that its content becomes as specific as it may be, even in any further circumstance c in which the utterance may be reconsidered. For example, "I" is construed in such a way that, once the reference of "I" in "I wear glasses" has been assigned in a context k, the resulting content being *<that x wears glasses>*, the content-in-context of the pronoun "I" is as specific as it may be, no matter the circumstances $c_j...c_i$ in which the original utterance may get reconsidered.

I take the Indexicality Assumption to hold at least for *pure* indexicals such as "I", "today", "tomorrow", i.e. for indexicals whose referent is fixed independently of the speaker's intentions. It may be pointed out that some "spurious" indexicals such as "here" or "now" do not comply with the Indexicality Assumption: in fact, since the referent of these expressions is partly determined by the speaker's intentions, fixing the referent of e.g. "here" as being location x in context k may not make the expression's content fully determined, for in other circumstances the speaker could have uttered "here" while referring to a wider location x'. Though I agree with this remark, I believe that the Indexicality Assumption still stands, for the following reasons: (i) the fact that the Indexicality Assumption doesn't apply in cases like "here" and "now" may be due to contingent facts about use of these words, which "interfere" with the way the term is designed to work semantically. It should not be forgotten that the Indexicality Assumption is not only a descriptive claim, but also a normative claim, telling us what indexical expressions are *designed* to do—rather than what they

merely actually do; (ii) moreover, it seems to me indisputable that, independently of putative exceptions like "here" or "now", the Indexicality Assumption captures the workings of at least some expressions which are widely recognised as indexicals and indeed as *pure* ones, like "I". So it seems that the Indexicality Assumption really captures a core feature of indexicality.

As opposed to what happens with (pure) indexicals, one might point out that the contextual completion of semantically under-determined sentences does not make their content-in-context determinate "once and for all". Let me illustrate: If a semantically under-determined expression e gets supplemented with extra content in a certain context k, it is not obvious that the underspecification of the expression is resolved for all the circumstances c_j … c_i in which the original utterance could be reconsidered. For instance, suppose that the content of "It's raining" is supplemented in context k_1 so as to result in *<it's raining in Paris>*. In some other circumstance c_2, the following question could arise: is it raining in Paris if it is raining only in the center of the city? In circumstance c_3, one could wonder: is it raining in Paris if it is raining over many scattered spots of the town? I envision that the content *<that it's raining in Paris>* that "It's raining" expresses in k_1 (*via* contextual enrichment) may turn out as not responding to the demands of greater precision that may arise in circumstances c_2 and c_3. That is, the content expressed by "It's raining" in k_1, *modulo* contextual supplementation, does not make the sentence's content determinate for all other circumstances c_2, c_3...c_n in which a higher standard of precision is required. After all, isn't it the case that the speaker who uttered "It's raining" in k_1 might have uttered the same sentence in other situations, meaning a wide range of contents—some of them richer in details, some other poorer—across these different circumstances? In a sense, then, given a certain contextual enrichment in context k, the semantic under-determinacy of an expression could be resolved in a certain circumstance c_k, but it could be "reiterated" in other circumstances c_j … c_i. Let me call this the *Reiteration Effect*. It is on the basis of these considerations that I believe that the following is true of semantically under determined sentences:

[Reiteration Effect] Once a semantically under-determined expression e has been completed in the circumstances c_k (relative to a context k), e may not express a determinate content in all circumstances of evaluation.[12]

Let the premise number (3) of our *reductio* be a statement of the *[Reiteration Effect]* claim. This ultimately results in the following argument, here presented schematically:

(1) A term like "green" is indexical;

(2) *[The Indexicality Assumption]* Indexicals (are designed to) work in such a way that, once the reference of an indexical term i is fixed in a context k, then i expresses a determinate content, in every circumstance of evaluation.

(3) *[Reiteration Effect]* Once a semantically under-determined expression e has been completed in the circumstances c_k (relative to a context k), e may not express a determinate content in all circumstances of evaluation;

---

[12] It may be pointed out that the Reiteration Effect is intrinsically problematic. Contextual enrichment delivers a content which may provide insufficient material for stating the truth-conditions of an utterance, because new ways of enriching that very content may arise. Ultimately then, no instance of enrichment can produce an evaluable content, because that content could be enriched indefinitely more and more. Let me clarify that this is an expected result: the fact that under-determinacy can be reiterated is precisely what makes semantic under-determinacy such a big problem for language as a means of communication. However, the Reiteration Effect is a *potential* problem, which doesn't actually threaten language users in concrete situations. This is because *the intentions of the speaker can set an upper limit to the possible enrichments an utterance's content can undergo*. This upper limit I take to be established by the speakers' intentions, which in turn depend on the interests, aims, expectations that are operative in the conversation. The fact that speakers' intentions can contain the impact of the Reiteration Effect is admittedly not a way of resolving the problem, but merely a way of taming it for the sake of successful communication.

(4) Suppose the indexicality of "green" were resolved in context k, so that "green" expressed the property *<being green on the outer surface>*;

(5) The under-specificity of "green" may have been resolved in the circumstances of evaluation c_k (relative to context k), but not in other circumstances of evaluation c_j... c_i.

(6) But this contradicts the *[Indexicality Assumption]*;

(7) Hence, (1) is false: a term like "green" is not an indexical.

As one can appreciate, the thrust of the argument is given by the fact that the defectiveness of semantically under-determined expressions gives rise to what I have called a *Reiteration Effect*: no matter how much one makes the predicate "being green" more precise *via* (alleged) indexicality resolution in a context k; in other circumstances of evaluation, it may continue to be indeterminate which property the predicate expresses. This "reiteration" of the predicate's under-specificity goes against the *[Indexicality Assumption]*, which (as far as I understand) captures a necessary trait of indexicality, i.e. the fact that the indeterminacy proper of an indexical term ceases—or at least it is expected to cease—for all circumstances of evaluation, once its referent has been fixed in a context[13]. This licenses the inference to (7), i.e. the negation of premise (1), to the effect that "green" is not an indexical.

As helpful as inference (1)-(7) might be, it doesn't take much to see that it is good only as far as it goes. For, even though, thanks to (1)-(7), the proponent of semantic under-determinacy can make sense of the distinction between semantically under-determined and indexical expressions, (1)-(7) only points at an *effect* of semantic under-determinacy (the *Reiteration Effect*), not at its *cause*. In other words, the argument set forth in (1)-(7) merely displays a *manifestation* of what semantic under-determinacy is, and not its source. But, as I stated at the beginning of **section 6**, my aim is to provide an answer to the *Demarcation Problem* by distinguishing semantic under-determinacy from other phenomena on the basis of what *gives rise* to it—and not on the basis of what effects it is associated with. The reason is that it seems that such a way of confronting the *Demarcation Problem* would provide a deeper response—to the extent that distinguishing phenomena by their causes seems more satisfactory from a theoretical point of view than distinguishing them on the basis of the effects or implications they are associated with. In what follows, I shall take a step in the direction of saying *in virtue of what* a semantically under-determined sentence is such: this will ultimately allow me to give a principled answer to the *Demarcation Problem*—as well as to display other significant theoretical benefits.

7. Under-articulated Sentences

One could start by drawing the following picture: our utterances often employ sentences that are *under-articulated*. For example, common speakers do not generally bother fully to articulate the salient location in their weather reports. This may be for all sorts of reasons: not articulating the

---

[13] It could be pointed out that even indexicals such as "here", "now" and demonstratives such as "this" and "that" are subject to the Reiteration Effect. This is undeniably true. Suppose that the referent of "here" in "It rains here" is fixed in context k, so that the content of the utterance in k is *<that it's raining in Paris>*. In different circumstances c_j, it could be asked whether that utterance of "It's raining here" could be true if it rains also in the suburban area surrounding the city of Paris. The k-relative content *<that it's raining in Paris>* doesn't help to establish this. This yields that "It's raining here" may turn out as under-determined even after indexicality resolution, if considered in different circumstances. I take it that this phenomenon doesn't falsify the *[Indexicality Assumption]*: independently of whether *uses* of indexicals are subject to under-determinacy, it still remains true that indexicals can be regarded as being *designed* so as to receive an absolutely determinate content for every circumstance. Moreover, it is sufficient that the *[Indexicality Assumption]* is made true by one example of indexical term. The expression "I" happens to be particularly suitable to this end, perhaps because its semantics is that of a "pure" indexical (see Kaplan 1989, Perry 2001). Thanks to Annalisa Coliva for discussion on this point.

location may save time, or articulating it would sound too unusual and would raise the suspicion in the hearer that the speaker intends to communicate something more. Think for example about Alan and Beth, who are walking together down the Rambla. At some point, rain begins to fall, but Beth doesn't notice. Alan then informs her by uttering "It's raining". No further articulation is supposedly required in order for Beth to understand what Alan said, as well as for her to learn whether it is true or false. She just needs to check out whether drops of rain are falling from the sky above her. Communication through the use of an under-articulated sentence has been successful in this case. Had Alan uttered "It's raining in Barcelona", Beth would probably have found it odder; she could have wondered whether Alan was picking up any discussion that they may have had previously, or she could think that, since Alan explicitly mentioned Barcelona without there being a reason, something more is at stake than simply the plain remark that it's raining: maybe Alan is trying to stress that *in Barcelona* it's raining, as opposed as what is happening in some other place. All sorts of hypotheses could be elicited in Beth's mind by Alan's remark to the effect that it's raining in Barcelona; on the other hand, the simple remark that it's raining would surely not get such a variety of reactions, and Alan would indeed be justified in choosing to utter "It's raining" in order to achieve his communicative aim.

There need be no calculation of the costs and benefits of what an utterance of a more elaborated sentence may imply. Speakers may just choose to be brief, if their audience already knows what they are talking about. Take for example Alan, who is in a bio-chemistry lab with Beth, choosing samples of green leaves in order to start analysing them. At some point, Alan comes across one of Pia's painted leaves. He recognises it as a Japanese maple leaf, and therefore concludes that it must be red underneath the paint. Showing the leaf to Beth, he utters "This leaf is not green. You'd better throw it in the garbage". Provided that Beth knows what Alan is talking about, no further articulation is needed in order for communication to succeed.

In both examples, speakers employ a sentence which leaves some element unspoken: in the first case, the location of the rain; in the second case, the way of being green of the leaf. Given what remains unspoken, the sentence's meaning doesn't determine the truth-conditions of its utterances, for the simple reason that *it doesn't articulate enough linguistic material* to that end. What remains unarticulated is some element that, had it been articulated, would have allowed the sentence to express the (intended) truth-condition of its utterance in some more exhaustive way. So, at least when we use sentences like "The leaves are green" or "It's raining", we may say that the meaning of these sentences doesn't give the truth-conditions of their utterances, for the sentence articulates too little material.

As the reader may have already realised, the picture I am adopting is not, at least so far, too different from that put forth by Quine and Sellars, as a way to deal with the issue of incomplete definite descriptions. i.e. definite descriptions which are not literally uniquely satisfied, but that nevertheless are used and understood as though they were uniquely satisfied (e.g. "the table", "the horse", "the house"). Quine says:

> Everyday use of descriptions […] is often elliptical, essential parts of the condition "...x..." being left understood; thus we may say simply "the yellow house" […] when what is to be understood is rather "the yellow house in the third block of Lee Street, Tulsa" (Quine 1940: p. 146)

Similarly Sellars, considering the sentence "The table is large" writes that

> a given utterance of it is elliptical and states what would be nonelliptically stated, for example, by "The table over here is large". (Sellars 1954: 200).

What does happen in those cases in which we use definite descriptions, as Quine and Sellars claim, "elliptically"? Here the account is in sync with the idea I have been outlining so far: when in our conversational setting enough elements are shared and, therefore, taken for granted by the conversational participants, speakers tend to be brief. Quine and Sellars' talk of "elliptical use" of a sentence thus captures the idea that speakers tend to be brief, and leave things unspoken. However, I believe that the expression "elliptical use", though effective for rhetorical purposes, should be pin-down more rigorously. Firstly, I believe that one should try to be clearer about what exactly is "elliptical", in an elliptical use. Usually, linguistic objects and not acts are elliptical. So the expression "elliptical use" is correct only insofar as it is a shorthand for "use of an elliptical linguistic expression". Secondly, ellipsis is a very precise phenomenon, which may be regarded as differing in significant respects from the notion we are currently trying to capture. Elliptical expressions typically "leave out" material which has already been articulated (in speech or writing). The sentences at focus here count as at least *un*-typical cases of ellipsis, in that they leave out elements which need not be antecedently articulated, but merely need be already salient "in the minds" of the conversation participants, so to speak (the rain location, the respect in which the leaf is green). For this reason, I would prefer not to use the term "ellipsis" to indicate this phenomenon, and I would opt for the more generic "under-articulation". My proposal is that of re-framing the idea of "elliptical use" in terms of use of under-articulated expressions.

At this point, though, the notion of "under-articulation" stands in need of a rigorous characterisation; because of course, there is little sense in the claim that sentences are under-articulated *tout court*—indeed, there is little sense in saying that, for any relevant predicate φ, an object o is under-φed *tout court*. Take for example the sentence "This country is under-developed": one can only evaluate the utterance if there is a comparison class, hence only if the country in question is less developed with respect to another country (or group of countries); the conditions for the possession of the property of under-development must include a comparison class. The same holds for under-articulation. In order for a sentence to be under-articulated, it must be under-articulated *with respect to something else*: in the present case, I surmise, another sentence.

I propose to regard the under-articulation of a certain sentence s as depending on a triadic relation between s, the (intended) truth-conditions of its utterance u, and a sentence s'. A sentence s will be under-articulated iff, when an utterance u of s is performed, another sentence s' could have been uttered, which articulates more linguistic material and has the same (intended) truth-conditions as u. For example if, talking about Paris, Eva utters "It's raining" while meaning <*that it's raining in Paris*>, we may say that the sentence that she utters is under-articulated iff Eva, in the same context, could have used the sentence "It's raining in Paris", which articulates more linguistic material and expresses the same intended content, and hence truth-conditions, of her actual utterance of "It's raining" as performed while talking about the weather in Paris.

Now the definition, as it is, may sound as a triviality: *every sentence is trivially under-articulated* in this sense, because for any sentence s, a more articulated sentence s' with the same truth-conditions as *any* utterance u of s could have been uttered. For example, "It's raining" is trivially under-articulated to the extent that one could have uttered "It's raining and it's raining", or "It's raining and 2+2=4" or "Oh my gosh, it's raining". All these sentences are more articulated than the simple "It's raining" and express the same truth-conditions of any utterance of this sentence: the first sentence does that in virtue of being the result of the conjunction of "It's raining" with itself; the second sentence does that in virtue of being the conjunction of "It's raining" with a necessary truth; and the third sentence does that in virtue of the fact that "Oh my God" has no truth-conditional role at all, and thus leaves the sentence's truth-conditions unaltered.

The definition of under-articulation I have just given may therefore appear too broad, and the notion uninteresting. The problem prompts me to call attention on a detail of the definition of under-articulation just provided; appreciating this detail will make for a principled way of ruling out uninteresting cases of under-articulation such as those listed above, as opposed to interesting cases.

What the definition states is that, in order for a sentence *s* to be under-articulated, *the more articulated sentence* s' *should be such that it linguistically articulates some truth-conditional aspects that were left out by the type-sentence* s, *even though they were part of the intended truth-conditions of the particular utterance, u, currently considered.* In other words, the more articulated sentence *s'* should *add some previously unexpressed truth-conditional information* with respect to the truth-conditions of the type-sentence *s*, in conformity with what the intended truth-conditions of the particular *utterance u* under consideration are.

This entails, for example, that the sentence "It's raining and it's raining" doesn't count as an interesting articulation of e.g. Eva's utterance of "It's raining", performed so as to mean *<that it's raining in Paris>*. *That* utterance is true iff it's raining in Paris; therefore, articulating Eva's originally uttered sentence "It's raining" entails putting into words the unexpressed propositional element *<in Paris>*. Conjoining a sentence with itself adds no truth-conditional information that is not *already* contained in the original type-sentence; it doesn't put into words what was left unspoken by the original utterance. Similarly, "Oh my god, it's raining" doesn't count as an interesting articulation of "It's raining", because the expression "Oh my God" has no truth-conditional import, and therefore leaves the truth-conditions of the original sentence unchanged.

Different considerations hold for "It's raining and 2+2=4": strictly speaking, the sentence qualifies as an articulation of "It's raining" in my sense, because conjoining it with "2+2=4" *does* add some extra truth-conditional information to the truth-conditions of the original sentence. The problem is: is there any utterance *u* of *s* which could have these truth-conditions? Or, better: is there any speaker who could *intend* an utterance of "It's raining" to be true iff it's raining and 2+2=4? This doesn't seem in principle inconceivable, however my intuition is that a context in which an utterance *u* of "It's raining" were true iff it's raining and 2+2=4 is rather far-fetched. Generally, utterances of "It's raining" are such that they leave out some piece of truth-conditional information concerning specific aspects of the raining-event described, as e.g. the location of raining, the substance involved in the raining (water, frogs...) etc. What is left out is usually a potential adjunct for the sentence, such that it could be articulated as: "It's raining *in Paris*" or "It's raining *frogs*". From a cognitive point of view, dropped adjuncts like "in Paris" or "frogs" are easily available to the hearer who wishes to identify the more articulated sentence *s'* that the speaker would have used in other circumstances. From these considerations it follows that, even though "It's raining and 2+2=4" is allowed as an articulation of "It's raining", a speaker who expected from a hearer to identify the former sentence would be "demanding too much" from his audience's inferential abilities (for similar concerns on the accessibility of the contents of under-articulated sentences, see Stanley 2002).

To sum up, in this section I have drawn attention to the phenomenon of *under-articulation*: it is a fact about our use of language that, when enough information is shared between a speaker and her audience, the speaker may choose to leave part of what she means unspoken. In such a case, the sentence that she uses is, in my view, under-articulated. A sentence *s* is under-articulated iff, for an utterance *u* of *s*, a more articulated sentence *s'* could have been used in the place of *s* to express the same content as that intended by the speaker who performed *u*.

8. Under-articulation as the Source of Semantic Under-determinacy

Having introduced and characterised the notion of under-articulation, we may now move to the key stage of the chapter. In this section, I will defend the claim that the semantic under-determinacy of sentences like "It's raining", "The leaves are green", "Jill is ready" etc. is explainable in terms (and only in terms) of under-articulation. This claim will prove helpful in several respects, as we will see in what follows. First of all, though, we need to relate back to the question that was asked in **section 6**, namely "What is the source of semantic under-determinacy?". This question I deem worth asking because it seems that a deeper understanding of semantic under-determinacy could be gained by

looking at what phenomenon could lie at the origin of it. As a way of answering this source-related question, I will therefore consider the following **Answer 3**:

Answer 3
[The Under-Articulation Claim] The meaning of a sentence *s* under-determines the content and truth conditions of its utterance *u*, because and only because *s* is under-articulated.

After having drawn attention on **Answer 3**, I wish to emphasise the fact that the semantic under-determinacy of sentences like "It's raining", "Jill is ready", etc. can be explained by employing the notion of under-articulation only. Take for example the sentence "It's raining", of which it was said that it is (at least most of the times) semantically under-determined. One may argue that the meaning of this sentence under-determines the truth-conditions of its utterances, solely for the reason that it doesn't put into words all that would make its utterances true or false. For instance, suppose Eva utters "It's raining" while meaning *<that it's raining in Paris>*. The reason why the sentence that she uses is under-determined may be just that it is under-articulated with respect to a longer sentence, like "It's raining in Paris", which could have been used in order to express that content. The same goes for "The leaves are green": one could explain the fact that the sentence is under-determined by appealing to the idea that the speaker who, by uttering it, meant *<that the leaves are green on the outer surface>*, could simply have used the more extended sentence "The leaves are green on the outer surface" in order to express that same content and truth-conditions. In general, my claim is that it is possible to explain the fact that a certain sentence *s* is semantically under-determined by appealing exclusively to the idea that *s* is under-articulated.
The causal-explanatory claim made by **Answer 3** is not to be conflated with the following "Emergentist" claim:

[Emergentist Claim] Every case of under-determinacy is a case of under-articulation.

The *[Emergentist Claim]* is too weak for my purposes, for, although it states something true in my perspective, namely that all cases of under-determinacy are also cases of under-articulation, it doesn't capture the *causal-explanatory* relation between under-articulation and under-determinacy. That this relation obtains is important, because it marks the fact that it is not by mere chance that every case of semantic under-determinacy is a case of under-articulation: there is a connection between the two phenomena which could be described as causal in some sense, to the extent that under-articulation, i.e. lack of linguistic material, gives rise to semantic under-determinacy, i.e. failure to provide the (intended) content in context and truth-conditions. Most importantly, though, the Under-articulation Claim should not be conflated with what may be called a "Reductive" Claim:

[Reductive Claim] Under-determinacy obtains if, and only if, under-articulation obtains.

This claim, besides failing to capture the causal link between under-articulation and under-determinacy, makes too strong a statement concerning the relation between the set of what is under-determined and the set of what is under-articulated. On the one hand, the left-to-right direction of the biconditional is true, for it just states that every case of under-determinacy is a case of under-articulation: as I just said, I maintain that, though I endorse also a causal statement. On the other hand, the right-to-left direction of the biconditional I think doesn't hold, for not every case of under-articulation is also a case of under-determinacy of meaning. Even if, in (probably) the majority of cases, under-articulated sentences do not express any determinate truth-condition, there may well be cases of under-articulation in which there is *no intuition as to semantic under-determinacy*. Let me clarify this point by making some examples.

An example is given by indexical expressions. Indexical expressions could be seen as under-articulated, in that instead of them, the speaker could have uttered a longer and more detailed definite description, or a proper name. For example, "I" is clearly under-articulated in this sense: suppose I utter "I am tired". The sentence is under-articulated to the extent that I could have used the sentence "The speaker of this utterance is tired" as well as "Delia is tired" in order to express the same truth-conditions as the original utterance. Leaving aside the semantic under-determinacy of "tired", the problem is that it doesn't seem that "I" suffers from semantic under-determinacy: the expression invariably has the same, fully determined character in every context and, moreover, once its reference is fixed in context, no residual under-specification is left, neither in the circumstances of evaluation of the context, nor in any other circumstance of evaluation. After all, as it was underlined in **section 6.2**, "I" responds to the *[Indexicality Assumption]*, which contradicts the claim about the *[Reiteration Effect]*. So, from the case of indexicals we may learn that not every case of under-articulation is a case of under-determinacy.

A further example is offered by Recanati's (2002a) weatherman case: suppose that rain has become very rare and rain detectors have been placed all over the territory. A monitoring station collects the records of every detecting device, and an alarm rings whenever rain is being detected. When the weatherman in charge of the monitoring room hears the alarm, he can utter "It's raining". His utterance may seem under-articulated, since he could have uttered "It's raining somewhere" or "It's raining in the area". However, in this case there is no intuition of semantic under-determinacy, since the sentence is true iff an event of raining is taking place *tout court* (in the intended region of space); that is, the sentence expresses an evaluable proposition even though no specific place is provided for the rain. So here again we seem to have a case of under-articulation to which no under-determinacy corresponds.

To sum up the content of this section: I have suggested that it is possible to explain every case of semantic under-determinacy as a case of linguistic under-articulation, along the lines of **Answer 3**. Whenever a speaker utters a sentence *s* whose semantic content under-determines what the speaker means with that utterance (and the related truth-conditions), it is possible to explain this phenomenon with the idea that the sentence *s* at issue just is under-articulated with respect to another sentence *s'* which the speaker could have uttered in its place to say what he meant with his actual utterance. In other words, **Answer 3** could be adopted as an account of the source of semantic under-determinacy. **Answer 3** should not be conflated with the *[Emergentist Claim]*, which just states that every case of under-determinacy is a case of under-articulation, but it doesn't capture the causal connection between the two; nor should it be conflated with the *[Reductionist Claim]*, which implies a commitment to the idea that every case of under-articulation is also a case of semantic under-determinacy—something which is not generally the case.

9. The Explanatory Advantages of Under-articulation

Let us survey the advantages that an under-articulation account would offer. First of all, a comparative advantage: recall the meaning indeterminacy view associated with **Answer 1** and the indexicality view associated with **Answer 2**: one can say that, if compared with these two competing views, an answer based on the idea that semantic under-determinacy is caused by under-articulation like **Answer 3** has the following benefits: first of all, it allows to maintain that meaning under-determines truth-conditions, but not because *meaning itself* is indeterminate; in fact, the meaning of sentences may be perfectly determinate; it's only that we use "too few words" to state what we mean. Secondly, the view allows to say that meaning under-determines truth-conditions, but such a failure is due to the simple fact that the sentence contains too few words to spell out the intended truth-conditions of its utterance, and not to the fact that it contains any indexical element beyond the acknowledged set. Other non-comparative advantages are the following:

(1) A Response to the Demarcation Problem. The under-articulation view helps the theorist to set apart semantic under-determinacy from other phenomena of semantic defectiveness, thus providing an answer to the *Demarcation Problem*. As already said, the under-articulation view (as stated in **Answer 3**) is not about a definition of semantic under-determinacy, but rather it is a view about what is sufficient as a *cause/explanation* of semantic under-determinacy. Consideration of what suffices to cause/explain semantic under-determinacy allows us to identify a feature which distinguishes semantic under-determinacy from other cases of semantic defectiveness. This feature consists in under-determinacy being *explainable in terms of under-articulation only*. This suffices in order to account for semantic under-determinacy as opposed to those phenomena like ambiguity, vagueness, indexicality, for it seems that all these phenomena are *not* explainable solely in terms of under-articulation. Indeed all of these phenomena seem to have their source in something *more* than simple shortage of linguistic material. For instance, it is not simply because of under-articulation that ambiguous words are such: ambiguous words are such because the same lexical item can be associated with different semantic properties. Nor is under-articulation the (sole) cause of vagueness: the root of vagueness seems to lie in the fact that the meanings of the words we use ("bald", "tall", "red") are not perfectly defined as to their application. Finally and most importantly, indexicality too isn't exhaustively explained by resorting to just under-articulation: indexical expressions are such not just because they substitute longer expressions, such as proper names or definite descriptions ("I" stands in place of "the speaker", "here" stands in place of "the location of the speaking"), but also because words can be used as devices for reference as a function of context; that is, language users can exploit meaning as a way of obtaining different contents in different contexts. In contrast, it seems that semantic under-determinacy is sufficiently or exhaustively explained by means of the under-articulation hypothesis only. Separate consideration will be reserved for ellipsis, which will be the subject of the next section.

(2) Explaining the Reiteration Effect. Another important advantage of the Under-Articulation View is that it offers the resources to explain, at a very basic explanatory level, the Reiteration Effect. We saw in **section 6.2** that, thanks to an argument like (1)-(7), we were able to distinguish the semantic results associated with semantic under-determinacy from those associated with indexicality by appealing to what I called a *Reiteration Effect*. The Reiteration Effect tells us that semantic under-determinacy may not be resolved "in all the possible circumstances of evaluation" because the same sentence s could have been used in different circumstances $c_1, c_2, c_3 \dots c_n$ to express different, more specific or richer, contents. This contrasted with the *Indexicality Assumption*, which stated that, once an indexical element has been assigned a reference in a context of utterance, it is not supposed to generate further under-specificity in circumstances of evaluation other than those related to the context of utterance.

If we adopt an account in terms of under-articulation, the explanation for why semantic under-determinacy may "reiterate" across different circumstance is simple: a single sentence s gives rise to the reiteration of semantic under-determinacy to the extent that it may under-articulate many sentences $s', s'', s''' \dots s^n$, each more articulated than the other, each of which could have been uttered in a different circumstance $c_1, c_2, c_3 \dots c_n$, to express each a different—more or less specific—content. So for example, one can say that the sentence

(4) The leaves are green

may give rise to a reiteration of semantic under-determinacy across different circumstances because it may be under-articulated with respect to the sentence

(4a) The leaves are green on the outer surface

but also with respect to the longer sentences

(4b) The leaves are green on the outer surface under normal light;
(4c) The leaves are green on the outer surface under normal light at 1 pm.

whose contents may have been intended by the speaker with an utterance of (4) in a range of different circumstances $c_1, c_2, c_3 \ldots c_n$.

(3) Explaining Context Shifting Cases. Last, but not the least, the under-articulation account explains the phenomenon that authors in the debate have considered as most indicative of semantic under-determinacy: the phenomenon of contextual shifts of truth-value exhibited by "Context Shifting" cases (see **section 4**). Remember the Pia case: while talking to the photographer friend, Pia utters "The leaves are green" and says something true; while talking to the botanist friend, Pia utters "The leaves are green" and says something false. Why this shift in truth-value? The explanation, according to the Under-articulation proponent, is that the sentence Pia uses is under-articulated with respect to two different sentences that she could have uttered in each distinct occasion, which would have semantically expressed two different contents and truth-conditions. The first occurrence of the sentence is under-articulated with respect to a sentence like "The leaves are green on the outer surface", which is true in the context of the conversation with the photographer friend; the second occurrence of the sentence is under-articulated with respect to a sentence like "The leaves are naturally green", which is false in the context of the conversation with the botanist friend. (Note that I am *not* claiming that these two more articulated sentences would have made the truth-conditions of Pia's utterances fully determinate. In order to claim this, one would have to ascertain how fine-grained the proposition Pia *intended* to express is. But this is not the enterprise I am now pursuing. Here my purpose is simply to give a *conceptualisation* of the phenomenon of semantic under-determinacy with respect to Pia's use of "The leaves are green"— not to *solve* or "dissolve" this under-determinacy).

To sum up, the Under-Articulation view has more than one advantage. Firstly, it is free of some uncomfortable assumptions, such as the assumption of meaning indeterminacy. Secondly, it allows to respond to the *Demarcation Problem* by tracing an intelligible distinction between semantic under-determinacy and indexicality (as well as ambiguity, vagueness, etc.), the distinction being that *semantic under-determinacy is explainable in terms of under-articulation only, while ambiguity, vagueness, indexicality are not*. The Under-Articulation view also allows to explain, at a very basic explanatory level, the so called "Reiteration Effect". Finally, it gives a basic and simple explanation of "Context Shifting" cases.

10. The Ellipsis Objection
At this point, the opponent of the under-articulation view could grant all the arguments so far developed but advance the following reserve: given all I have said about semantic under-determinacy, the result is just that I am claiming that semantic under-determinacy is explainable in terms of a form of *ellipsis*. If this is true, however, then I haven't gone very far: all I have been doing is, after all, domesticating semantic under-determinacy, because I have related it to an already familiar phenomenon.
A response to this objection requires making a distinction. Under-articulation surely has some features in common with ellipsis: both of them are about linguistic material that is "left out". Anyhow, there are important (even though subtle enough) differences between the two phenomena. Ellipsis I take to be a purely *syntactic* phenomenon, which is invoked by linguists in order to make sense of the grammaticality of certain sentences. On the other hand, I take under-articulation to be a

semantic-conceptual phenomenon. To see the difference between the two, suppose the following dialogue takes place:

(6)   Alan: "Who baked the cake?"
        Beth: "Jill"

The response given by Beth is not a sentence at all, however in this context it sounds as a good way of answering the question posed by Alan. The explanation for this is that the non-sentential item uttered by Beth is *elliptical*: at the overt level, the verb phrase has been deleted, however it is retrievable in one of the previously articulated sentences. In this case, it's easy to find the previously articulated item: it's the interrogative sentence "Who baked the cake?". The verb phrase "baked the cake" is superficially absent in Beth's utterance of "Jill", even though an interpretation of this utterance requires to see the non-sentential item as completed exactly by such a verb phrase, at least at the covert level. By considering this example, we are already able to identify one of the essential marks of ellipsis as a syntactic phenomenon: for a certain sentence (or non-sentential linguistic item) to be syntactically elliptical, it has to leave out some material which is *already present in some previously articulated sentence*. In other words, when an expression is elliptical, it is because the *linguistic context* can provide the linguistic material that the expression leaves unspoken.[14]

This feature is not shared by cases of under-articulation. For an under-articulated sentence could be uttered even though no previous sentence has been uttered before: that is, under-articulated sentences can occur in a discourse-initial position. Suppose Beth and Alan have been discussing extensively the day before whether Jill is prepared for her Spanish exam. The day after, Beth and Alan meet again. Beth asks Alan "So is Jill really ready?". As one can see, as long as both Beth and Alan already know what the topic of the conversation is, the sentence "Jill is ready" can be felicitously used to start their conversation (consider that what has been articulated in the previous conversation may not provide the syntactic material for completing the sentence). Under-articulated sentences can then occur discourse-initially if by "discourse-initially" we mean at the beginning of a discourse or conversation within an already shared background of presuppositions. No previous *sentence* needs to be articulated within the same conversation: it is sufficient that the speakers presuppose certain *propositions*.

Given these considerations, we may legitimately conclude that explaining semantic under-determinacy in terms of sentential under-articulation doesn't amount to explaining it in terms of ellipsis; what I have been doing is not domesticating semantic under-determinacy, i.e. relating it to an already familiar phenomenon in natural language. Quite the contrary: by explaining semantic under-determinacy in terms of under-articulation I have committed to the idea that semantic under-determinacy arises merely because of the way speakers *employ* language (they tend to be brief whenever enough is presupposed in the conversation), and not because of any mechanism grounded in language in any deeper semantic/syntactic sense.

---

[14]   Neale (2004) traces the distinction between syntactic ellipsis and under-articulation (which he calls "utterance ellipsis") in the following way: syntactic ellipsis is a genuinely *linguistic phenomenon* while under-articulation is merely a *way of using language*. Syntactic ellipsis is the deletion of one or more sentential components, which can be easily recovered in the material antecedently articulated, thereby allowing the reconstruction of the non-elliptical sentence via a "copy-paste" process. Uttering a syntactically elliptical sentence implies engaging in an elliptical use. However, the reverse doesn't necessarily happen. This is because one can have mere *ellipsis of use*: for example, one might *use elliptically* the sentence "The table is large", in the sense that one could have used some longer sentence, However, the sentence "The table is large" is not syntactically elliptical at all. Under-articulation is here the equivalent of Neale's ellipsis of use. As Neale admits, his insights are essentially in the footsteps of Quine (1940) and Sellars (1954).

Conclusion

In this chapter, I have introduced and discussed semantic under-determinacy, i.e. the idea to the effect that the meaning (linguistic content) of a sentence *s* fails to determine the truth-conditions of any utterance of *s*. After having historically framed the problem, I have structured the chapter as follows:

- I have drawn attention onto what I have called the *Demarcation Problem*, which may be summarised by the question "Is there a feature that sets apart semantic under-determinacy from other familiar phenomena of semantic defectiveness like ambiguity, indexicality, vagueness, ellipsis?".

- I have presented and discussed the Context-Shifting Argument (CSA), which is the argument usually advocated by proponents of semantic under-determinacy in order to establish their thesis. I have then emphasised these authors' efforts to distinguish semantic under-determinacy from other phenomena of under-determinacy such as ambiguity, indexicality, vagueness, ellipsis, thus pursuing what I called a *no-domestication policy*. I have stressed that the no-domestication policy doesn't give us the means to deal with the *Demarcation Problem*;

- As a first step towards answering the *Demarcation Problem*, I have asked the question "What is the source of semantic under-determinacy?". According to **Answer 1**, the meaning of a sentence doesn't determine its utterances' truth-conditions because meaning *itself* is indeterminate. According to **Answer 2**, meaning is not indeterminate in itself, but it is still too poor or unspecific in order to provide the *content-in-context* of utterances;

- I have attacked **Answer 1** by arguing that, unless further qualifications are provided, it is incompatible with the intuitive and valuable distinction between *semantic* correctness and *pragmatic* correctness. It seems that, whenever a speaker uses a semantically under-determined sentence, she is using that sentence for what it means. The under-determinacy is not caused by an unstable meaning, but by an incomplete, stable meaning;

- I have also criticised **Answer 2**, on the account that its formulation doesn't allow to distinguish between semantic under-determinacy and indexicality. Moreover, an argument like the one in (1)-(7), purported to set the two phenomena apart, only draws attention on an effect of semantic under-determinacy, which I call the *Reiteration Effect*. What we need is, however, a *cause* (explanation) for such an effect;

- I have suggested to think of semantic under-determinacy as caused (explained) by *sentential under-articulation*. It is just a matter of fact that, in appropriate circumstances, speakers do not care about using sentences that fully articulate the truth-conditions of their utterances (e.g. "It's raining"). Under-articulation thus obtains in all those cases in which, when an utterance *u* of a sentence *s* is performed, a sentence *s'* could have been used instead of *s*, such that (i) *s'* articulates more linguistic material than *s*; (ii) it has the same (intended) truth-conditions as the particular utterance *u* of *s* currently under consideration;

- The main explanatory advantage of the under-articulation approach is that it provides an answer to the *Demarcation Problem*: what sets apart semantic under-determinacy from other phenomena of semantic defectiveness such as ambiguity, vagueness and indexicality is that semantic under-determinacy is *explainable in terms of under-articulation only*. The main difference between semantic under-determinacy and ellipsis is that felicitous use of elliptical sentences needs that some material is previously articulated, while this is not the case for semantically under-determined expressions;

- Moreover, the under-articulation account (a) explains the "Reiteration Effect" at a very basic explanatory level, by saying that under-determinacy reiterates to the extent that a certain sentence *s* is under-articulated with respect to a range of more articulated sentence $s', s'' \ldots s^n$ which could have been used in different circumstances $c_1, c_2 \ldots c_n$ to articulate each a

different—more or less rich—content; (b) explains the context-shifting effect: a sentence s is true in context c_1 and false in context c_2 because it under-articulates two different sentences s_1 and s_2 such that, if uttered each in one of the two contexts, would have yielded different contents and truth-conditions;

Chapter 2

Semantic Under-determinacy, Comprehension and Radical Contextualism

In the previous chapter, the issue of semantic under-determinacy was introduced and explored in its most fundamental aspects. The chapter's central aim was to answer the *Demarcation Problem*, i.e. to positively characterise semantic under-determinacy as opposed to ambiguity, indexicality, vagueness, and ellipsis. I argued that semantic under-determinacy is *explainable solely in terms of under-articulation* (where under-articulation interestingly differs from ellipsis). This feature sets apart semantic under-determinacy from other phenomena of semantic defectiveness, such as ambiguity, vagueness and most importantly, indexicality. All these phenomena appear as not explainable by an appeal to linguistic under-articulation only—all of them have their source in something different than simple shortage of linguistic material. For instance, it is not simply because of under-articulation that ambiguous words are such; nor is under-articulation the (sole) cause of vagueness; finally and most importantly, indexicality too isn't exhaustively explained by resorting to just under-articulation. In contrast, it seems that semantic under-determinacy is sufficiently or exhaustively explained by means of the under-articulation hypothesis only.

In **Chapter 1**, I maintained that a response to the *Demarcation Problem* was called for as a preliminary step towards a treatment of semantic under-determinacy. Once this preliminary step has been taken, moving to the next stage requires turning to a perhaps even more challenging project, which is that of explaining *how, in general, communication is possible despite semantic under-determinacy*. The rest of the thesis is devoted to discussing answers previously given to this question, with the purpose to assess them and to finally put forward an independent answer. In the present chapter I will explore one of the most radical answers that can be provided for this problem. Other answers will be assessed in **Chapters 3** and **4;** an independent proposal will be advanced in **Chapter 5**.

1. The Comprehension Problem and the Contextualist Answer

Observing that language is subject to semantic under-determinacy cannot but make the following question pressing: how is it that, despite the under-determinacy exhibited by what is encoded in sentences, utterances manage to be understood and evaluated, and communication manages to turn out (at least generally) successful? The problem is stated by considering the point of view of the *potential listener* of one of the sentences that have been surveyed so far. What is to be explained is the general mechanism by which the listener of an utterance of "It's raining", "Jill is ready", "The leaves are green" etc. manages to understand what the utterer means. Since the issue specifically concerns the *comprehension* of under-determined utterances, it could be formulated as follows:

[The Comprehension Problem] What is the general mechanism by which the listener of a semantically under-determined sentence *s* understands what the speaker means by uttering *s*?

A variety of answers is compatible with the Comprehension Problem. In the following chapters, I will consider different versions of an approach according to which the right answer to the Comprehension Problem is that context determines some important aspects of the semantics of expressions. This is a broadly conceived "Contextualist" answer, and it could be captured as follows:

[The Contextualist Answer] The context in which a sentence *s* is used determines aspects of the semantics of the utterance *u* of *s*.

There is more than one way in which the Contextualist Answer may be read, for there is more than one semantic aspect which could be affected by elements of the context. (1) The first aspect which could be affected by context is the *linguistic meaning* of an expression. With regard to this semantic aspect, the Contextualist Answer is to be read as stating that there is a relation of determination between the linguistic meaning of the expressions that we use and the particular context in which we happen to use them; that is to say, context affects meaning itself. (2) The second aspect that could be amenable to contextual effects is the *content-in-context of an utterance*. The resulting reading of the Contextualist Answer will be that context operates on sentences endowed with a stable conventional meaning, which nevertheless are assigned a different content as context changes. There are two ways of understanding this reading: (2.1) The expression's meaning, although stable, determines a *sub-propositional* item which is subsequently enriched in context by pragmatic processes that operate independently of the semantics of language. So, for example, according to some theorists, the sentence "The leaves are green" determines a sub-propositional item, which gets enriched pragmatically (i.e. in a linguistically unconstrained way) in order to deliver a truth-evaluable item. In this conception, the relation between the context of utterance and the semantics of the expression consists in the fact that context freely *enriches* the expression's sub-propositional content. (2.2) The expression's meaning is stable, but it consists of a *function* from a context of utterance to a content. So for example, according to some theorists, the word "know" denotes a function from a context of utterance to a knowledge-relation (see DeRose 1992, Cohen 2000). Other expressions, e.g. indexicals, clearly refer to different objects in different contexts. Their meaning is thus such that it denotes a function from the context of utterance to a content: For example, "I" means invariantly "the speaker" but, if uttered by Alice it refers to Alice, while if uttered by Beth it refers to Beth, etc. In this "functional" conception of the semantics of certain expressions, the role of the context is that of providing a value for the function, and hence the relation between context and semantic value will be a relation of "*functional saturation*". (3) The last possible reading of the Contextualist Answer focusses on that semantic aspect of an expression known as its *extension*. The extension of a predicate is a set of individuals: so, for example, the extension of "bachelor" is the set of all the individuals that are bachelors. The extension of a sentence is a truth-value: for example, the extension of "Charles is a bachelor" is the truth if Charles is a bachelor and falsity otherwise. As soon as extensions become in focus, the Contextualist Answer can be interpreted as saying that there is a relation of determination between the extension of the expressions that we use and the particular *context* in which we happen to use them, i.e. that extension is sensitive to the variation of certain contextual features.

In the present chapter, I will briefly deal with reading (1), which I shall dismiss as incompatible with a truth-conditional semantics. I will then move on to reading (2.1), which I shall consider more extensively, and against which I will argue that it brings to a merely partial view of the way speakers understand language in context. I will turn to reading (2.2) in **Chapter 3** and to reading (3) in **Chapter 4**.

2. Extreme Contextualism

Reading (1) of the Contextualist Answer clearly espouses a view of semantic under-determinacy as a feature of the *meaning* of words and sentences. For, if what is semantically under-determined is meaning itself, then in order for utterances of sentences to be true or false, there will have to be some interaction between *meaning* and context, in the sense that context shall have a role in determining what words and sentences linguistically mean. The resulting adaptation of the Contextualist Answer corresponds to what I shall call The Extreme Contextualist Answer

[The Extreme Contextualist Answer] Relevantly different uses of an expression *e* can determine relevantly different meanings for *e*.

However, as it stands, the Extreme Contextualist Answer does need some qualification. Having it that meaning depends on use in this way cannot mean allowing every use of a word to mean whatever the speaker wishes to mean. Embracing the Extreme Contextualist Answer with no further qualification is compatible with an endorsement of the claim that meaning is not subject to any form of constraint, such that it allows building up a stable semantics for the language in question. What interests me here is, however, the possibility to maintain that words and complex expressions in a language have a stable semantics, even though this semantics may under-determine the content of the utterances made by speakers in concrete contexts. I will, therefore, disregard the Extreme Contextualist Answer as not enough interesting for the purposes of an investigation about the semantics of a language.

3. Radical Contextualism

Let us then move on to reading (2.1) of the Contextualist Answer, which is less extreme for it has it that, notwithstanding their having a stable linguistic meaning, utterances of some expressions—as such—do not strictly speaking "say" anything. Context then freely (i.e. independently of any linguistic constraints) operates on the content of these sentences. The corresponding adaptation of reading (2.1) could be called a "Radical Contextualist Answer":

[Radical Contextualist Answer] An expression *e* has a determinate linguistic meaning, however free (i.e. linguistically unconstrained) pragmatic processes determine "what is said" by the speaker with an utterance of *e* in relevantly different contexts.

In light of this claim, one may want to distinguish between two senses of "what is said". On the one hand, an expression says something in virtue of the semantic properties it possesses by convention: this is the sense of "what is said" coined by Grice (1967/1989). Proponents of the Radical Contextualist Answer obviously have to grant that expressions in English say *something* in this Gricean sense. On the other hand, an expression says something if its content determines a complete proposition, i.e. an evaluable content, in context. As we have seen in **Chapter 1**, theorists such as Travis (1985/2008) and Recanati (1989) contend that expressions such as "The leaves are green", "It's raining" say *nothing* in this second sense, i.e. they do not express any truth-condition in context.

According to Radical Contextualism, what the sentence says in the traditional, Gricean sense is regarded as indispensable, for it provides minimal semantic directions. However, it is often hardly of relevance for the reconstruction of utterance truth-conditions. The meaning that is encoded in sentences is nothing but a *schema or template*, which provides very general and unspecific semantic information, but ultimately doesn't express any truth-evaluable content. It is at this point that the role of pragmatic processes becomes valuable, for it helps the hearer to determine "what is said" by the speaker in the second sense—i.e. to determine a proposition which is evaluable in context. The processes that serve this task are usually called processes of "free enrichment", to mark the fact that they supplement conventional semantic content with additional propositional material in a way that is "free" from linguistic constraints. Processes of free enrichment will be the subject of the next section.

4. Processes of Free Enrichment

One of the mostly advertised benefits of Radical Contextualism is that it seems to deal pretty well with the problem of semantic under-determinacy. Sperber and Wilson (1986/1995, 1998), Carston

(2002) and Recanati, (2004) are among the Radical Contextualists that put a stress on the problem and take a substantial step in order to overcome it. This step consists in positing *pragmatic effects* affecting what sentences say in context. But what is the nature of such processes, and how do they set themselves apart from other pragmatic processes that are commonly thought of as perfectly integrated in natural language semantics?

We said that, according to Radical Contextualism, what makes "what is said" by a speaker with an utterance *u* of a sentence *s* richer and more specific than "what is said" by *s* in virtue of meaning alone is the effect of *pragmatic processes* which exploit information drawn from the context of utterance. At first glance, this is reminiscent of what in formal semantics has been described as the phenomenon of *indexicality,* i.e. the semantic exploiting of the context of an utterance aimed at picking out the reference of certain expressions—like "I", "here", "this". Since Bar-Hillel (1954), Montague (1974), Davidson (1967a) Kaplan (1989), Evans (1982) and Perry (1993), indexicality has been recognised as a legitimate module of the semantics of natural language: it just happens that the content of expressions like "I", "you", "now", "that" needs to be "anchored" to the extralinguistic context in which these expressions are used, for various reasons that pertain to the efficiency and economy of the language tool for the sake of successful communication. So a resort to the context of utterance in order to fix the content (namely the denotations) of some expressions was already provided by natural language semanticists before Radical Contextualism showed up. What is then the originality of the Radical Contextualist suggestion?

Here the Radical Contextualist project importantly diverges from the project of modelling indexicality within the semantics of natural languages. For the mechanisms underlying indexicality are fundamentally *constrained by linguistic forms*. For each instance of indexicality there is an element, which is articulated at the linguistic (lexical or syntactic) level, which calls for saturation, i.e. to which a referent has to be assigned. The point of Radical Contextualism is precisely to claim that there are processes which exploit the context which, however, go beyond the paradigm of indexicality, in that they do not respond to linguistic constraints but are triggered by *purely pragmatic demands*. As Recanati writes:

> Various pragmatic processes come into play in the very determination of what is said; not merely saturation—the contextual assignment of values to indexicals and free variables in the logical form of the utterance—but also free enrichment and other processes which are *not linguistically triggered but are pragmatic through and through.* (Recanati 2004: 21, my emphasis)

A similar remark is expressed by Robyn Carston:

> [T]here is no level of linguistic representation of [… a] sentence in which there are variables (or silent indexicals or empty constituent slots) which indicate that the contextual values must be assigned in order to determine the full truth-conditional content. […] An adequate account of how these meaning-constituents become part of the proposition expressed by the utterance […] is formulated entirely in terms of pragmatic mechanisms. (Carston 2004: 818)

All these non-linguistically-driven processes can be equated with what Radical Contextualists call the phenomenon of *free enrichment*. Free enrichment is an umbrella-term, that covers a variety of processes in which context freely (i.e. free from linguistic constraints) affects the meaning of sentences so as to finally return some richer, more specific, or in any case more fleshed-out content for the utterances of such sentences. It should be noted that the commitment to free pragmatic processes is what fundamentally sets apart reading (2.1) from the claim contained in reading (2.2): while in reading (2.2) meaning is equated with a function from context to content, whose role is that

of constraining the effects of context on what a certain utterance says, in reading (2.1) the role of context and of the related pragmatic effects is thought of as free from linguistic strictures. In light of this, one may wonder what the propositional elements are that are supplied by such a range of free, unconstrained processes, as well as how free enrichment functions precisely.

[Unarticulated Constituents]. The first category of elements that are provided by processes of free-enrichment is composed by so-called *unarticulated constituents*. The notion of unarticulated constituent was first introduced by Perry, who made it vivid through the following example:

> It is a rainy Saturday morning in Palo Alto. I have plans for tennis. But my younger son looks out the window and says, "It is raining". I go back to sleep.
> What my son said was true, because it was raining in Palo Alto. There were all sorts of places where it wasn't raining: it doesn't just rain or not, it rains in some places while not raining in others. In order to assign a truth value to my son's statement, as I just did, I needed a place. But no component of his statement stood for a place. The verb "raining" supplies the relation *rains (t, p)*—a dyadic relation between times and places, as we have just noted. The tensed auxiliary "is" supplies a time, the time at which the statement was made. "It" doesn't supply anything, but is just syntactic filler. So Palo Alto is a constituent of the content of my son's remark, which no component of his statement designated; it is an unarticulated constituent. (Perry 1986: 138)

What Perry in this passage calls "unarticulated constituent" is a propositional component that does not correspond to any articulated component of the sentence, either at the superficial, overt level or at the deep, covert one. That is, at no level of the syntax of "It's raining", whether superficial or deep[15], is there an element that stands for the location—Palo Alto—which that particular statement of "It's raining" concerns. (This of course doesn't rule out that the verb itself "raining" encodes, at least at the lexical level, a relation between times and places *rains (t, p)*).

Perry's notion of unarticulated constituent is key for Radical Contextualists, who subscribe to the view that context supplies propositional components in linguistically unconstrained ways. For these theorists, unarticulated constituents are brought into the proposition in ways that do not respond to linguistic strictures: in the case reported above, for example, the locational element *<Palo Alto>* is supplied by the hearer just in order for the utterance of "It's raining" to be evaluable and as informative and relevant as possible given the purposes of the current communicative exchange.

By means of which processes are unarticulated constituents incorporated into the proposition expressed? In what follows, I will focus on two accounts of how this might be accomplished. The first one, due to Sperber and Wilson (1986/1995, 1998), puts a stress on the *cognitive* processes by which a certain content gets pragmatically enriched, which the authors characterise as following a "path of least effort". The second account, due to Recanati (2002a) (but traceable to McConnell-Ginet (1982)), emphasises the way pragmatics may trigger important transformations at the syntactic level.

[The Path of Least Effort Approach] In Sperber and Wilson's framework, sentences like "It's raining" correspond to mental representations whose structure admits of a certain range of expansions or, as they say, "developments". Supplying an unarticulated constituent ultimately means developing this mentally represented "conceptual structure" in the ways that are more appropriate, given the context of utterance. The psychological process which underlies this

---

[15]   Perry is very explicit about the fact that the unarticulatedness of the constituent pertains both the sentence's superficial and *deep* syntax. As he maintains, "we do not need first to find an expression, hidden in the ʹdeep structureʹ or somewhere else, and then do the semantics of the statement augmented by the hidden expression." (Perry 1986: 143).

"development of conceptual structure" responds to what Sperber and Wilson call a "principle of optimal relevance". What this principle entails in practice is that the subject who witnesses an utterance of, e.g., "It's raining", engages in a search for the most cognitively fruitful interpretation of it, at the least cognitive effort. The strategy that Sperber and Wilson presume to be followed by competent interpreters is the so-called "path of least effort": basically, the hearer chooses the interpretation that most quickly comes to mind, which best satisfies the current expectations of relevance. Here is how the authors describe the process:

> The hearer takes the conceptual structure constructed by linguistic decoding; following a path of least effort, he enriches this at the explicit level […], until the resulting interpretation meets his expectations of relevance; at which point, he stops. (Sperber and Wilson 1998: 192)

In Sperber and Wilson's view, the "path of least effort" consists in a non-demonstrative, sub-personally run *inference*, which leads the hearer to identify the proposition that best satisfies the requirement for optimality of relevance. Let us consider, as a basic example, the following exchange:

(1) A: "What's the weather in Paris?"
    B: "It's raining"

Now B's utterance is such that it doesn't articulate the place where the rain is falling. Therefore, B's utterance can appear to A as a proper answer to her question only if A engages in an extra piece of reasoning (even a tacit or sub-personal one) in which she exploits the decoded linguistic meaning of B's utterance as a premise (or input) and, through a finite number of steps, finally arrives at the propositional content that best fits with A's expectations of relevance and B's presumed intentions. The reasoning could be reconstructed thus:

(1) B uttered the sentence "It's raining";
(2) If B's utterance is about Paris, then it is optimally relevant.
(3) Since A made a question about the weather in Paris, the answer must be about Paris, too;
(4) If B's utterance had not been about Paris, B would have let A know explicitly;
(5) B's utterance is about Paris;
(6) B's utterance is optimally relevant;
(7) Hence, the content expressed by the utterance is that it is raining *in Paris*.

One can see how, through a finite number of inferential passages, B arrives at an enriched reconstruction of the content of A's utterance. In this process, the conceptual structure of the predicate "raining" is thought of as being enriched by turning it into the more complex conceptual structure *raining_in_x*, where *x* occupies a "free slot" in the conceptual structure, dedicated to locations. If the variable *x* is filled in with the locational element *<Paris>*, the final upshot is that the content of A's utterance is understood in context as being *that it's raining in Paris*, rather than simply *that it's raining*. Sperber and Wilson are committed to the view that this process responds, at the cognitive level, to a "least-effort" strategy, which demands to the hearer that he retrieves the most relevant interpretation at the lowest cognitive cost.

[The Variadic Function Approach] The second account of how unarticulated constituents are supplied is more concerned with the issue of how we can reconcile a compositional, hence syntactically controlled, account of free-enrichment with the idea that the process in question is purely pragmatically driven. The solution consists in positing that *pragmatic processes are able to modify the syntactic structure of expressions*. Recanati (2002a) takes it that unarticulated

constituents are introduced in the proposition via a modification of the *logical form* of the sentence's occurrence. In order to achieve this result, he posits an operator named *Variadic Function* which, provided the presence of a syntactic or even contextually relevant predicate-modifier, takes as input a certain predicate *R(x)* and gives as output a predicate whose logical form is increased (but also *de*creased, if appropriate) by one argument place, as e.g. the two-place predicate *R(x, y).* So suppose that the verb in question is "to rain". Recanati holds that, in absence of any syntactic modifier like "in Paris" or of any contextually salient location, the verb expresses the monadic predicate *rain(t)*, designating a property of times. However, if an explicit modifier is made to precede the predicate, or if the predicate is used while making a specific location salient, the logical form of the predicate acquires an extra argument-place. As an effect of the variadic function, it goes from being the one-place relation to times *rain(t)* to being a two-place relation *rain(t, l)* between times and locations. In addition to this, given that either the sentence or the context already contains the element suitable to saturate the newly acquired argument-place, the process gives as output not just the "bare" expanded logical form *rain(t, l)* but the very same form with the novel argument-place *already filled-in*. As Recanati writes:

> [… W]e can say 'It is raining' without providing a location for the rain, whether linguistically or even contextually. [...] But if we do provide a location, either through the adjunction of a prepositional phrase or by purely contextual means, *we thereby generate a new relation, in which there is an empty slot, an argument role which the location fills*. Since there is an argument place for the location in the output relation (though not in the input relation), the location finds its way into the proposition and coheres with the other constituents. (Recanati 2002a: 322)

Consider another example: "Jill is ready"; suppose that an utterance of this sentence is intended by the speaker to be true if, and only if, Jill is ready *to take the exam*. The interpretation of the word "ready" then requires that the predicate *being ready* is turned via variadic function into a different one, namely the predicate *being ready_for*, which contains an argument-place to be filled in with the propositional element *<the exam>*. The ultimate result is that the content-in-context that the utterance of "Jill is ready" expresses is the richer content *that Jill is ready to take the exam*. Of course, had the relevant purpose of Jill been different, what the occurrence of "ready" would have said in that context would have been different, too.

To sum up, I have introduced the notion of "unarticulated constituent", which is a component of the proposition expressed by an utterance *u* of a sentence *s*, to which no articulated (syntactic, lexical) elements in the sentence *s* correspond (neither at the level of superficial or at the level of deep syntax). Sperber and Wilson regard unarticulated constituents as the result of a process of "development of conceptual structure" which follows a strategy from the least cognitive effort. Recanati is more worried about how pragmatic demands may modify the syntactic structure of expressions, so as to preserve compositionality, and therefore resorts to an explanation in terms of what he calls "Variadic functions".

[Ad-hoc Meanings]. The second category of elements coming as a product of free-enrichment includes what we may call *ad-hoc meanings:* an *ad-hoc* meaning is constructed out of some core semantic traits, by adding traits that are purely contextually derived. Work on *ad-hoc* meanings is intimately related to work on the formation of concepts, which was carried out mainly by cognitive psychologist Larry Barsalou (1983): very sketchily, Barsalou realised that humans are able to form categories which organise the world according to very contingent, temporary and particular purposes: for instance, the category "things to eat on a diet"; the category of "noises coming from the neighbours' house"; the category of "things that can be used as hammers", etc. In so doing, subjects create concepts that are "ad-hoc", i.e. which fit contingent situations and the related tasks

to be carried out, such as eating low-calories food, recognising annoying noises, nailing objects to the wall. Inspired by Barsalou's work, Sperber & Wilson (1998: 194-96) and Carston (2002: 143-52) have characterised instances of free-enrichment as cases of "ad-hoc" meaning formation, i.e. of cases of formation of meanings which suit contextual purposes. Let us see how.

In their view, an utterance of a word invokes a *concept*, which goes into the mental representation associated with the sentence uttered. In the case of semantically under-determined expressions, the concept invoked and, hence, the whole sentence's representation, is too poor (or too rich, in some special cases) in order for it to capture what the utterer is saying. This means that the concept at issue needs to be adjusted with respect to what is relevant in the context—it needs to be made "ad-hoc"; the word's meaning, the proposition expressed, as well as the utterance's truth-conditions are then affected accordingly. The difference between *ad-hoc* meaning formation and supplementation of unarticulated constituents is that, while in the case of unarticulated constituents a development of conceptual form is required in order to make room for an extra constituent, in the case of *ad-hoc* meanings no room for extra-constituents needs to be prepared, for enrichment involves the *semantic* traits of a word, rather than its structural ones.

The mental process corresponding to *ad-hoc* concept formation is, again, an *inference* aimed at making utterances in their particular context of utterance *optimally relevant*, i.e. conveying as much cognitively useful information as possible, at the least cognitive cost. Consider, for example, the following exchange:

> (2) A: "How about going to the movies?"
>   B: "I'm tired"

Here the answer counts as relevant only if B's being tired is a reason for not going to the cinema. So the hearer needs to run an inferential reasoning aimed at picking the right sense of "tired". Once such sense is retrieved, and the right truth conditions for (10b) are worked out, the hearer can therefore move on to inferring what B is implicating, namely that B is not feeling like going to the movies (the derivation of implicatures was something that we covered in **Chapter 1**). The inference gone through by A might be as follows:

1. B has uttered "I am tired";
2. In the context of this utterance, B's utterance is optimally relevant only if B is saying that he is tired to an extent that prevents him from going to the movies;
3. If what B is saying was not that he is tired to an extent that prevents him from going to the movies, B would have let me know;
4. What B is saying is that he is tired to an extent that prevents him from going to the movies;
5. B's utterance is optimally relevant;
6. Hence, the content expressed by the utterance is that B is tired *to an extent that prevents him from going to the movie*s.

What is important in the relevance-theoretic framework is that contextual enrichment modifies the meaning of the sentence by modifying the very *concept* associated with a word by a mental representation. Thus, in view of the contextual data available, the hearer will presumably recognise that the concept corresponding to the word "tired" is not the simple concept TIRED, but a more fine-grained concept TIRED*. Thus, according to relevance theorists, the mental/ conceptual representation of the meaning-in-context expressed by "I am tired" will be

(10b') tired*(B)

Corresponding to a more specific proposition like

(10b″) B is tired to an extent that prevents him from going to the movies;

Once the hearer has identified (10b″) as what the speaker means, she can proceed to calculate what the speaker is implicating—namely, that he doesn't want to go to the cinema.

To sum up, so far I have illustrated two ways in which a process of pragmatic enrichment may help with the problem of semantic under-determinacy: providing *unarticulated constituents* and creating *ad-hoc meanings*. The underlying strategy is the same in both cases: have sentential content-in-context enriched so that the expressed content amounts to a complete, evaluable proposition, while at the same time eluding the constraints posed by indexicality.

5. "Deep" and "Shallow" Semantic Under-determinacy

Before moving on, let me put forward some qualifications aimed at restricting the scope of the present investigation, so as to avoid future misunderstaning as much as possible. I wish to trace a distinction between *genuine* semantic under-determinacy and a kind of under-determinacy which is merely pragmatic. So far we have characterised semantic under-determinacy as a failure on the part of linguistic meaning (in context) to determine a truth-condition, which implies that a gap obtains between what the sentence literally says (in context) and what an utterance of that same sentence means—according to the speaker's intentions.

In a first range of cases, the reason for this gap between what is strictly-speaking said and what is meant by the utterance is that the meaning of the sentence at issue delivers an altogether incomplete proposition. The gap therefore obtains between a *non-propositional* semantic item on the one hand and a *genuinely propositional* item on the other. The cases of "The leaves are green", "Jill is ready" and "It's raining" are paradigmatic in this sense. Neither of these sentences expresses a truth-condition in ordinary cases, and thus a gap opens between what these sentences say and what a speaker may manage to say by using these sentences in a context. This is what I will call "deep" semantic under-determinacy.

In a second range of cases, even though the meaning of the sentence delivers a complete proposition with complete truth-conditions, the proposition literally expressed is not what speakers usually express by using that sentence. A gap opens between a *propositional* literal sentence meaning on the one hand and an equally *propositional*, even though more "appropriate" utterance content. Take for example an utterance of the sentence "Jack and Jill got married". Strictly speaking, this sentence is true if, and only if, Jack got married and Jill got married. This is however not what is generally meant and understood: in general, when this sentence is uttered in a particular context, the intended (and grasped) content is that Jack and Jill got married *to each other* (see Carston 1988). Or, take an utterance of "It will take time to get there". Literally, the sentence says that it will take a certain amount of time to get to some destination, and it is true iff it will take a certain amount of time to get to that destination. However, this would be a triviality, not worth stating in normal circumstances. The content that is usually intended and that usually gets grasped is instead that it will take *more time than expected* in order to get to the destination (see Carston 2002). Finally, consider an utterance of "I have had breakfast". Strictly and literally, the sentence is true iff the speaker had breakfast at least once in his life. However, if I utter such a sentence at ten o'clock, in response to a friend's offer of a sandwich, this is not the content that I am most likely to express. What I rather express, and what my hearer understands, is the richer proposition to the effect that I have had breakfast that very morning (see Recanati 2004).

As far as I can see, only cases of "deep" under-determinacy exemplify the right way of conceiving semantic under-determinacy for the purposes of the present work, i.e. as lack of determinacy of *meaning* (in context) with respect to *utterance truth-conditions*. Even though cases in the second group, too, exhibit some form of under-determinacy, I think that, in virtue of the fact that the sentences in question do express evaluable propositions, the under-determinacy at issue does not deserve to be called "semantic". It is probably a *pragmatic* under-determinacy, i.e. under-

determinacy of what would be usually appropriate to express in a certain occasion with an utterance of that sentence. I will therefore disregard cases belonging to the second group, and focus exclusively on the first group.

Radical Contextualism has it that processes of enrichment operate on both these groups of cases; however, it is in the case of "deeply" semantically under-determined sentences that the crucial role of pragmatics with respect to semantics becomes apparent. This is because processes of enrichment *need* to be carried out in order for utterances of these expressions to express a truth-evaluable content in the first place. As Carston says, semantically under-determined expressions

> require a pragmatic process of completion before they can be judged as true or false descriptions of a state of affairs. (Carston 1999: 105)

What is at stake, here, and what is attained with the help of pragmatics, is the possibility that these sentences express a truth-evaluable content, and not just a content which is relevant, informative, conversationally helpful. It's at this point that pragmatics appears, at least in the view of authors such as Carston and others, as truly indispensable for the semantics of these expressions, and that we can start talking about *Truth-Conditional Pragmatics*.

In other words, I believe that the sense of "semantic under-determinacy" relevant for the purposes of this research should be restricted to *deep* semantic under-determinacy, which is a failure on the part of linguistic meaning (in context) to determine utterance content as a consequence of its failure to determine *an evaluable content at all*. Cases in which linguistic meaning (in context) fails to determine utterance content which are due to the fact that the semantic content is a full-fledged proposition, though not the proposition which the speaker intends—which may be called "shallow" cases of semantic under-determinacy—will not be in the focus of the present research. The project I am after is eminently that of investigating how communication is possible whenever meaning under-determinacy is deep, i.e. it is such that it defeats evaluation.

6. Radical Contextualism and Comprehension

Let us now introduce the issue of comprehension. As previously stated, the [Comprehension Problem] can be summarised as the question: *how, in general, do hearers manage to understand utterances of semantically under-determined sentences?*

In what follows, I shall assume that comprehension is nothing more than knowledge of the truth-conditions of an utterance. Knowledge of truth-conditions doesn't mean merely knowledge of a T-sentence. There is more to understanding an utterance of "Snow is white" than simply being able to derive the sentence "'Snow is white' is true iff snow is white". I take it that a subject who competently derives and therefore comes to know the truth-conditions of an utterance of "Snow is white" also knows *how the world should be* in order for the utterance to be true.[16] This means that understanding an utterance involves not just "blindly" applying the rules of a formal semantic system, but also having this system interact with one's knowledge of how the world is, broadly speaking.

Furthermore, I take it as indisputable that, if there is no (determinate) truth-condition, there can be *no knowledge* of truth-conditions. Thus, if an utterance of "It's raining" expresses no (determinate) truth-condition, one can have no knowledge of this utterance's truth-condition, and the same applies for any utterance of "Jill is ready" and "The leaves are green". But we defined comprehension as knowledge of an utterance's truth-conditions. This entails that, if there is no (determinate) truth-condition expressed, there is no comprehension of the utterance at issue.

Now let us go back to Carston's point that semantically under-determined sentences need pragmatic enrichment in order to express a truth-condition, and therefore be evaluated as true or false. In light

---

[16]   For a debate on the notion of comprehension as related to a truth-conditional theory of meaning see Davidson (1967, 1976), Foster (1976), McDowell (1980), Rumfitt (2005), Longworth (2008).

of these considerations, the upshot is that, if no enrichment is carried out, then no truth-condition is expressed by an utterance, and no comprehension (i.e. knowledge of its truth-condition) is possible. But what kind of notion of "comprehension" is this? Radical Contextualists cannot be working with *any* notion of comprehension: what they seem to have in mind is a notion of comprehension that amounts to knowledge of the truth-conditions of an utterance of a sentence in the *particular context* in which the utterance is performed—given the speaker's intentions, the operative presuppositions, the salient objects in context, etc. Let us call this notion *RC-Comprehension.*

[RC-Comprehension] Knowledge of the truth-conditions of an utterance in the particular context in which it is performed.

With this notion of comprehension at hand it is certainly true that, if no enrichment can be carried out, then there can be no comprehension of an utterance of the sentence "It's raining", for there can be no knowledge of the truth-conditions that that utterance has in the particular context in which it is performed, with the particular intentions of the speaker and conversational setting that pertains to that very utterance.

Let us delve a little more into this issue. What appears as crucial for attaining RC-Comprehension of an utterance is, I take it, the exercise of the hearer's pragmatic-inferential capacities in order to figure out what the speaker means while using a certain linguistic expression in a particular context. The exercise of these pragmatic-inferential capacities is aimed not at linguistic de-coding, but at making the content of the speaker's utterance as "cognitively useful" as possible in particular contexts. Processes such as the interpretive inferences envisaged by Carston and Sperber & Wilson, and described in **section 4**, take as a premise the semantic content of the sentence and work out what the speaker meant by searching for the most cognitively useful interpretation at the lowest computational cost. This is accomplished by means of inferential, largely abductive strategies. So the role of meaning is confined to that of providing the premises for what are, after all, genuine inferential procedures directed at gaining comprehension of a certain linguistic act. As these authors themselves claim:

> While still assuming that the code model provides the framework for a general theory of communication, and hence for a theory of verbal communication, most pragmatists have described comprehension as an inferential process. (Sperber and Wilson 1986/1995: 12)

> The derivation of the proposition explicitly communicated is dependent on pragmatic inference, not merely in determining intended referents and intended senses of ambiguous expressions, but in supplying unarticulated constituents and adjusting encoded conceptual content (enriching and/or loosening it). (Carston 1999: 86)

Pragmatic-inferential processes aimed at comprehension of what the speaker is saying with a certain utterance have a crucial feature, which is so obvious that it goes almost unnoticed: The inferential abilities which take contextual information into consideration have to be exercised within *particular* contexts of utterance. It would be pointless to run these inferential abilities without being acquainted with a context, for the point of running them is exactly establishing what the speaker intends to say—something which requires being familiar with the context in which he is speaking. For example, it would be pointless for a hearer to exercise her pragmatic-inferential abilities in order to establish what a speaker is saying with an utterance of "It's raining" without resorting to the particular context in which the speaker is performing the utterances. This task requires the presence of a particular context which supplies the ingredients to be supplemented to the expressed content— for instance, the location of the rain.

From this it follows, quite obviously, that if no such particular context is accessible, the hearer cannot figure out the intended truth-conditions expressed by the utterance at interest and no RC-comprehension of that utterance can take place. This is because, if the hearer is not familiar with the context in which a certain utterance is performed, she cannot fruitfully exercise the pragmatic-inferential abilities which would enable her to know the truth conditions of the utterance in that particular context, and hence she won't attain RC-Comprehension.[17]

Let us summarise the results of the present section. The notion of comprehension that apparently lies behind Radical Contextualism was labelled as *RC-comprehension*, i.e. knowledge of the truth-conditions of an utterance in the particular context in which it is performed. Inferential-pragmatic abilities lie at the root of RC-Comprehension, and they most fruitfully unfold within particular contexts, for the the aim of such processes is determining what the speaker intends to say in that particular context of utterance. From this it follows that, if no particular context is accessible to the hearer, no inferential-pragmatic abilities can be successfully exercised, thus preventing the hearer to understand what the speaker means in the particular context in which he performs the utterance. Ultimately then, in absence of a particular context, the hearer can attain no RC-comprehension of an utterance (of a semantically under-determined expression).

7. Intuitive Comprehension

Whether or not RC-Comprehension is a viable notion (I think it is), there is nevertheless a sense of comprehension according to which we do understand an utterance of "It's raining" even if we have no access to the particular context in which the utterance is taking place—for example because we ignore which particular location the speaker is talking about. This is because, I claim, hearers are able to figure out the truth-conditions that any utterance of a sentence would have in any arbitrary context. I shall call this *Intuitive Comprehension*:

[Intuitive Comprehension] Knowledge of the truth-conditions of any utterance of a sentence in any arbitrary context.

Having an intuitive comprehension of an occurrence of e.g. "It's raining" implies that we know that any utterance of such a sentence is true in any context iff it's raining somewhere or other. Similarly, there is a sense of comprehension in which we perfectly understand an utterance of "Jill is ready" even if we don't know which particular activity is under discussion. Our intuitive understanding consists in the fact that, upon hearing an utterance of "Jill is ready", we thereby know that any utterance of this sentence is true in any context iff Jill is ready for some purpose or activity. In general, I shall say that intuitive comprehension is grasp of the truth-conditions that any arbitrary

---

[17] As an aside, it is also acknowledged by authors in the debate on the methodology of semantics of natural language that, in the comprehension of untutored speakers, *general* or *abstract* sentence properties and *particular*, *context-dependent*, utterance properties are going to be tightly intertwined. Untutored speakers' reports as to what a certain utterance "says" or "means" are thus going to be "contaminated" by considerations pertaining to the utterance's significance in a particular context. As Cappelen and Lepore observe: "the actual use of locutions like 'means that' (and 'said that' […]) reveals no 'pure' semantic intuitions about these uses. Intuitions about when such reports are true, it seems, are always cut with non-semantic material: often they are about what a speaker said, sometimes about what he meant, sometimes attempts to explain what was said, sometimes they provide lexical or syntactic or pragmatic information, almost invariably they are a mix of several of these" (Cappelen and Lepore 1997: 290). All this obviously raises question as to whether interrogation of reports from untutored speakers is good methodology for the sake of setting forth a semantic theory of a more or less big fragment of natural language. See e.g. Devitt (forthcoming) for the idea that intuitions of experts are more reliable than those of common speakers as indicators of the structure of a semantic competence.

utterance of a (semantically under-determined) sentence has *in any arbitrary context.*[18] [19]

Is there an independent motivation, besides the considerations pertaining to the present discussion, for taking the notion of intuitive comprehension seriously? As far as I can see, independently of the issue of semantic under-determinacy one could grant it that there is a sense of understanding on the basis of which we are able, independently of our comprehension of an utterance in a particular context, to know in general what it is for any arbitrary utterance of that sentence to be true in an arbitrary context. This difference between understanding a particular utterance and understanding any arbitrary utterance is apparent with respect to e.g. definite descriptions: independently of the referent that one can understand a particular utterance of "the teacher of Alex" to have in a particular context, one is able to understand any arbitrary utterance of "the teacher of Alex" as meaning "the unique x such that x is the teacher of Alex". Intuitive comprehension seems to have a place in our semantic intuitions even independently of the question of our comprehension of semantically under-determined expressions.

What would a Radical Contextualist say about Intuitive Comprehension? Nothing encouraging, I suspect. We already saw that, when no access to contextual information is possible, Radical Contextualism predicts, rather obviously, that there will be no RC-comprehension, for there will be no chance to run any pragmatic inference—which is required for that kind of understanding. However, as far as I can see, Radical Contextualism also predicts that there won't be any intuitive comprehension either. Let us see why.

Semantically under-determined sentences do not yield any truth-conditions for their utterances. Intuitive understanding requires access to a truth-conditional content, namely in the form of knowledge of the truth-conditions of any arbitrary utterance in any arbitrary context. If there is no particular context available, it will be pointless to run any pragmatic inferences in order to discover the content and truth-conditions intended by the speaker. So the hearer has to stay content with the sentence's conventional semantic content (in context). However, conventional semantic content doesn't give the truth-conditions of *any utterance whatsoever*. So the hearer who doesn't have an access to the utterance's context will not be in a position to know the truth-conditions of any arbitrary utterance in any arbitrary context, thereby being unable to intuitively understand the sentence's occurrence.

Suppose now that the hearer were not in any particular context, but that he were able to locate the utterance in a *default* context, in which the sentence expresses at least a *default* content. In this case, RC-Comprehension would still be possible. However, if even no default context were available, the hearer would be doomed to neither RC-understand nor intuitively understand what the semantically under-determined sentence says. Following the characterisation of intuitive understanding we have given, we get it that outside of a particular or of a default context, the hearer won't be able to know what the truth-conditions of any arbitrary utterance of that sentence in any arbitrary context are.

Let us make some examples. Take an utterance of "It's raining". If a particular context is available in which Tokyo is salient as a location for the rain, the utterance will be RC-understood as being

---

[18]   Some may equate *intuitive* comprehension with *linguistic* comprehension, i.e. comprehension of an utterance of a sentence which proceeds from one's competence about linguistic meaning. If linguistic competence enables one to know the truth-conditions that any utterance of "It's raining" expresses no matter the context in which is used, then I am happy to equate intuitive comprehension with linguistic comprehension. Note that, in any case, this context-independent kind of comprehension is to be distinguished from the deliverances of a formal semantic system. I take it that doing the compositional semantics of "It's raining" simply delivers a sentence like "'It's raining' is true iff it's raining". Being able to derive such a sentence is *not* the same as knowing, in a context-independent way, in what conditions *any utterance* of "It's raining" is true or false.

[19]   The intuitiveness here is due not only to the fact that such comprehension is conscious, i.e. that the intuition is a mental episode which is present to the subject's consciousness, but also to the fact that it involves interrogating one's competence as to utterance-interpretation. Note that here I am not using the term "intuition" in the non-epistemic sense in which some theorists used it, for example as "opinion" (Lewis 1983), "intellectual seeming" (Bealer 1998), or "inclination to believe" (Sosa 2000). Here intuition is an epistemic notion to the extent that it involves *knowledge* of the truth-conditions of an arbitrary utterance of a sentence in an arbitrary context.

true iff it's raining in Tokyo. If there is no such context, the hearer may still provide a *default* context for the utterance: for example, he may presume that most of the times, when one utters "It's raining" one means that it's raining in the place of the utterance. He may therefore RC-understand the utterance as being true iff it's raining in the place of the utterance—say Paris. However, if no such default context were available for any reason, there would be no "starting point" for any interpretive pragmatic inference. The hearer should stay content with the sentence's semantic content (in context), namely *<that it's raining>*. However, this doesn't provide the content of any utterance whatsoever. As a consequence, there would be no intuitive understanding either, for the sentence "It's raining" just semantically expresses the sub-propositional content *that it's raining*, which doesn't express any truth-conditional content in any context. Another example could be the following: take an utterance of "Sam cut the cake". If a particular context is available in which Sam is cutting the cake with a chainsaw, the utterance will be RC-understood as being true iff Sam cuts the cake with a chainsaw. Suppose, now, no context were available: comprehension of the sentence may still involve fixing some *default* features of the contexts in which utterances of "Sam cut the cake" are usually true, e.g. that Sam is using a knife. By providing a default context, RC-comprehension is guaranteed. But suppose that there was no reason to select such a default context. Absent any context, no pragmatic inference can get started in the context. The hearer has to stay content with the sentence's content-in-context *<that Sam cut the cake>*. This doesn't provide him with the content of any utterance whatsoever. The hearer then won't have any intuitive understanding, because the simple semantic content of "Sam cut the cake" doesn't express any truth-condition in any context.

To sum up, in this section I have tried to show how, in the Radical Contextualist framework, claims (1)-(4) entail (5).

(1) The meaning (in context) of semantically under-determined sentences doesn't determine any utterance truth-conditions;

(2) Intuitive understanding is knowledge of the truth-conditions of any arbitrary utterance of a sentence in any arbitrary context;

(3) If no context is available, no pragmatic inference can be run, and the hearer has to stay content with the meaning (in context) of the semantically under-determined sentence.

(4) If (1), then the meaning of a semantically under-determined sentence fails to determine the truth-conditions of any arbitrary utterance in any arbitrary context;

(5) If no context is available, there is no intuitive comprehension of a semantically under-determined sentence, because it fails to determine the truth-conditions of any arbitrary utterance in any arbitrary context.

8. Intuitive Comprehension out of a Context

In this section, I will challenge claim (5), according to which out of a particular context (and in absence of a default context) there can be no intuitive comprehension of an utterance of a semantically under-determined sentence. I will argue that language users do have such intuitive comprehension with respect to semantically under-determined sentences. A little experiment, call it *The Eavesdropper experiment*, is going to illustrate the kind of case I would like to make.

> *[The Eavesdropper Experiment]*
> Suppose you overhear someone shouting "Jill is ready" from the next room, but you do not know who uttered that and what state of affairs the utterance is intended to describe. Now suppose you try and answer the following two questions:
>
> (i) Do I know, independently of the context, in which case the utterance would be true?
> (ii) Do I know, independently of the context, whether the sentence is true or false?

The answer to (i) is plausibly going to be positive. First of all, you plausibly have understood what has been said by the *sentence*, as long as you sufficiently master a semantics for the English language, in the sense that you master a compositional, syntax-driven machinery, which would allow you to derive the biconditional "'Jill is ready' is true iff Jill is ready". But this is not all for intuitive understanding. Besides having such a "mechanically" derived comprehension of the sentence, you can also claim to know more: you know what has to be the case in the world in order for any utterance of that sentence to be true—namely, it has to be the case that Jill is ready *for some purpose, or activity*. This, I surmise, is because you plausibly know that correct use of any utterance of "X is ready" is such only if X is ready for some purpose. In virtue of your being familiar with what amounts as a correct use of the predicate "being ready", you may know that any arbitrary utterance of that sentence is going to be true in any context just in case Jill is ready *for some purpose* in the circumstances in which the utterance was issued.

Let us now turn to question (ii). The answer is plausibly going to be negative. For, you have no idea whether that sentence, uttered by the unknown person for unknown reasons, aiming to describe an unknown state of affairs, is true or false. In other words, you have no elements to know whether there is any state of affairs in which an utterance of that sentence is true.

The answers to the previous questions seem fairly uncontroversial. Are they uncontroversial for the Radical Contextualist, too? Let's consider first the answer to question (ii). Here the prediction of the Radical Contextualist parallels the intuitive one, and it would be surprising if it didn't. Independently of whether one is a Radical Contextualist, it is simply reasonable to suppose (and it would be crazy to deny) that, unless one is provided with the *relevant* data concerning the circumstances in which the utterance takes place, one is not going to be in a position to judge whether a overheard (or over-read) occurrence of a sentence like "Jill is ready" is true or false[20].

Let's now focus on the answer to question (i). Here is where the intuitive response diverges from the one that ensues from Radical Contextualism, for the Radical Contextualist would predict that the answer to question (i) is *negative*. Why would it be so? Because, if one takes seriously the claim that intuitive understanding is knowledge of the truth-conditions that arbitrary utterances of a certain sentence have, then hearers will be predicted not to intuitively understand utterances of semantically under-determined sentences out of a context, because these determine *no truth-condition* for their utterances. But but this just seems false. Even out of a context, a speaker knows what it takes to make utterances of "Jill is ready" true: namely, that Jill be ready *for something* (a purpose, an activity). The same holds for "It's raining": even in absence of any contextual information, the hearer knows that usually utterances of that sentence are true iff it's raining *in some place or other*[21]. Even in the case of some more weird-sounding sentence like "Sam cut the sand", there is an intuitive understanding of the conditions in which utterances of this sentence would be true, namely iff Sam cuts the sand *in some way or other*.

The Radical Contextualist may be skeptical about this way of describing intuitive comprehension. A critical observation that he could advance is the following: there is no such thing as an intuitive understanding of e.g. "Jill is ready". Speakers merely think they have this intuitive, non-context-specific understanding, but what is really going on is that they *implicitly provide a default or "made up" context* in which the utterance may be said true or false.

This kind of skepticism seems at least unwarranted: why should we think that there is no such thing as intuitive understanding? It seems reasonable that our understanding of semantically under-determined sentences out of their contexts is guided by considerations as general as possible, in absence of any particular contextual clues: these general considerations may well be semantic/conceptual in nature, i.e. they could very plausibly involve reflection on what counts as

---

[20] Of course, there are exceptions: e.g. math sentences, tautologies, contradictions.
[21] Here there is probably a restriction on the existential quantification; so, for example, what the hearer understands is not simply that it's raining somewhere, but that it's raining somewhere *in the area salient to the speaker*.

the correct use of a predicate like "to rain", "being ready" etc. On the contrary, it would seem unreasonable for us to reconstruct a default or rather a completely made-up context, if we had no good reasons to do that. So in this case, the Radical Contextualist owes us an explanation of how it can be the case that, *in general*, it's preferable to see comprehension out of a context as involving the reconstruction of a default or made-up context rather than the consideration of general facts about the correct application of the predicates employed by the speaker.

9. Comprehension and Conceptual Competence

As we have seen from the previous section, speakers can have intuitive comprehension out of a context, despite the negative prediction issued by the Radical Contextualist. I contend that the reason why there can be an intuitive comprehension of an utterance of "It's raining" or "Jill is ready" or "The leaves are green" or "Sam cut the sand" is that the speaker can rely on what we will call *conceptual constraints* over the notions of raining, of being ready, of being green and of cutting. What are conceptual constraints? Conceptual constraints are *ways of structuring and organising thought about certain specific matters.* We can think of a conceptual constraint as a rule that specifies what has to be the case in order for the use of a certain concept to be correct. For example, a conceptual constraint on correct uses of the concept TO RAIN has it that the concept is correctly applied only if it rains at a location. A conceptual constraint on uses of the concept of BEING READY is the one to the effect that an application of the concept of readiness is correct only if the readiness is related to some activity. A conceptual constraint on CUTTING is such that an application of this concept is correct only if the cutting is effected in some way, or with some method of cutting. A conceptual constraint on BEING GREEN is such that a use of the concept of being green is correct just in case the greenness is greenness in some respect. Conceptual constraints may be regarded as giving very general and fundamental directions about what counts as a correct application of a concept, which is to say that they give directions as to how thought and other cognitive tasks concerning specific matters should be conducted.

Is there any difference between conceptual constraints and *semantic* constraints? Of course, conceptual constraints are constraints on thought, while semantic constraints are constraints on meaning. Concepts are certainly related to meanings: knowing the meaning of a term, e.g. "rain", entails mastering the related concept RAIN. On the assumption that hearers of "It's raining", "Jill is ready", "The leaves are green", etc. are also competent language-users, conceptual constraints certainly *coincide* with semantic constraints for them. Note in passing that, in the present case, the notion of "semantics" can't be used any more to refer to a formal, syntax-driven compositional machinery. If we want conceptual constraints to be also semantic constraints, we have to broaden our notion of semantics: in this conception, semantics provides the truth-conditions for utterances of the sentences of a language *L*, but it does so in a way which is not only syntactically constrained, for it is also constrained by rules that govern *correct use* of words/concepts (see especially **section 6** of **Chapter 5** on this). The ensuing relation between semantics, syntax and the conceptual domain could be represented as follows:

Semantics (determines truth-conditions of utterances of sentences in *L*)	**responds to**	**Syntactic** constraints
		Conceptual (semantic) constraints

Though conceptual constraints *de facto* coincide with semantic ones for competent language-users, I will nevertheless tend to privilege talk about conceptual constraints. This is because I believe that there is a sense in which conceptual constraints are more *stable* than semantic constraints. Let me explain this claim. The greater stability of conceptual constraints can be seen by considering that mastering a concept doesn't necessarily entail that one knows the meaning of the term which, by hypothesis, is "associated" with that concept in a certain language. There are at least two cases in

which this obtains. The simplest (but trivial) case is the one in which one doesn't know the term to begin with: That is, one could be competent about the concept *F* but fail to know the term *e* which, in language *L*, is associated with *F*. More complex (and interesting) cases could be imagined, in which one is competent about concept *F*, knows the term *e* which is associated with *F* in *L*, but is not competent about the meaning of *e*. For instance, philologists, who presumably are competent about the ordinary concept of hospitality, may yet have false beliefs about the meaning of the word "xenia" in ancient greek, which (let us suppose) may have been associated with exactly the contemporary, ordinary concept of hospitality which they master. Not only this. The greater stability of conceptual competence as opposed to semantic competence also arises from considering cases in which a subject gains comprehension of a certain word *w* in a language *L* which, by hypothesis, is associated with concept *F*, thanks to precisely her mastery of *F*. For instance, the subject could observe the contexts in which *w* occurs (the sentences it is used in, the occasions in which it is employed), relate these contexts to those in which the concept *F* is correctly used, and hence justifiedly conclude that *w* has a certain meaning which is relevantly related with *F*.

What follows from these remarks is that concepts (and conceptual constraints) are not tied to language the way semantic constraints are. They are more general devices of thought-organisation, which can be possessed independently of one's semantic competence about an arbitrary language *L*. These considerations suggest that conceptual constraints are more stable than semantic constraints, and thus provide a motivation for adopting a view of the comprehension of semantically under-determined sentences which relates it with competence on conceptual (rather than purely semantic) constraints. It goes without saying that, as I already stressed, for all those hearers who are also competent language-users, conceptual constraints just *coincide* with semantic constraints. In what follows, though, I will nevertheless privilege talk in terms of conceptual constraints, and I only occasionally switch to talk of semantic constraints.

Mastering and applying conceptual constraints depends on whether a subject is a competent concept-user. Whether or not one is a competent concept-user depends on whether one is enough trained, i.e. on how experienced one is with uses of the concept at issue. Experience in this respect is acquired through the repeated use of concepts within a community, which implies gradually adjusting one's competence to that of the other members of the community, by proceeding through errors, corrections and reinforcement of correct uses. Being conceptually competent means complying to the rules of a shared cognitive system, i.e. it implies conforming the structure of one's thought and reasoning about certain matters to the one that the rest of the community shares and currently employs. We may say that conceptual constraints arise from the fact that subjects who share an environment, be it natural, cultural, social or linguistic, generally wind up sharing a *standard cognitive format* for thought on a range of different matters, whose function is that of favouring inter-subjective coordination, or subjective coordination with the world. In other words, conceptual constraints are ways of structuring and organising thought about certain matters which promote and facilitate coordination at various levels—the coordination of an individual with her environment, with other subjects and so on. As for the relation of conceptual constraints with language, it take it that conceptual constraints are shared by those speakers who are also competent language-users, in virtue of the fact that mastering the meaning of an expression entails mastering the concept which is also associated with it.

What is the relationship between conceptual constraints and utterance comprehension, in particular intuitive utterance comprehension of semantically under-determined sentences? The idea, in a nutshell is that *conceptual constraints help gain intuitive comprehension of utterances of semantically under-determined sentences*: Suppose that an utterance of "It's raining" is performed, and the hearer has no access to any feature of the related context. To the extent that the hearer is aware of the conceptual (semantic) constraint that to rain is to rain *at a location*, she is able to figure out the intuitive truth-conditions of a semantically under-determined occurrence of "It's raining", even out of a context, i.e. she is able to know that an utterance of that sentence is true in any

arbitrary context iff it's raining somewhere. Analogously, to the extent that the hearer is aware of the conceptual (semantic) constraint to the effect that being ready is invariably being ready *for some purpose/activity*, she is able to figure out the intuitive truth-conditions of a semantically under-determined occurrence of "Jill is ready" even out of a context, i.e. she is able to know that an utterance of this sentence is true in any arbitrary context iff Jill is ready for some activity. The same remarks apply for "green" and "to cut".

Note that, once conceptual constraints enter the scene, intuitive comprehension immediately stands out as *inferential but at the same time non-context-specific*. It is inferential, because inference is still required in order for the hearer to figure out that an occurrence of "It's raining" is true just in case it's raining somewhere (in the area salient to the speaker). The hearer who employs the conceptual constraints related to the concept RAIN may run an inference of the following form in order to obtain the intuitive truth-conditions of an utterance of "It's raining". Having acknowledged that the speaker uttered "It's raining", the hearer may continue as follows:

(1) Rain necessarily takes place at a location;
(2) All uses of "It's raining" must be true iff it's raining somewhere;
(3) There is no reason to suppose that this use of "It's raining" differs from the others;
(4) So this utterance is true iff it's raining somewhere.

This inference interestingly differs from the procedures envisaged by Carston and Sperber & Wilson (reviewed in **section 4**). First of all, those may be regarded as having the structure of *inferences to the best explanation*: they start with some linguistic inputs and employ heuristic strategies guided by the least effort principle, in order to recover the intended content. These inferences involve the exercise of rather complex cognitive abilities—such as the ability of "mind reading"—and notions—such as the notions of relevance, of information etc. Successfully running these inferences requires applying strategies that involve a fair amount of creativity on the part of the hearer, together with an inclination to explain, rather than mechanically decode. Nothing like this is the case with relation to the (1)-(4) inference above, which starts with a linguistic input, employs generalisations based on conceptual (semantic) constraints and eventually concludes—absent any reason to think otherwise—with a prediction about the content expressed. Clearly running this kind of inference requires much less creativity and is not purported to *explain* anything. This inference may be gone through in a mechanical way, since the application itself of conceptual (semantic) constraints is supposed to be a mechanical task, an operation which may be seen as belonging to a subject's deeply entrenched cognitive (and discursive) habits, in virtue of her sharing a standard cognitive and discursive format on these matters.

Intuitive understanding is also clearly non context-specific. No specific context is needed by the hearer in order to figure out the intuitive truth-conditions of an utterance of "Jill is ready" or "It's raining". This is precisely because intuitive understanding exploits features of the concepts READY and RAIN that are *cross-contextual*, i.e. that apply no matter the context, because they pertain to general features of these concepts (as well as of the related predicates) that are known in virtue of one's conceptual competence and one's sharing a standard format for thought and discourse on these matters with other members of the same natural, cultural, social, linguistic group.

10. The Significance of Conceptual Constraints

All these considerations cast a new light on Radical Contextualism and its notion of pragmatic enrichment. As we have seen, enrichment processes are *inferential*, *context-specific* processes. With the notion of conceptual constraints, however, we seem to have uncovered a more basic level of understanding, which is still inferential, yet definitely *non-context-specific*, in that it exploits cross-contextual information that is drawn from one's having one's thought (and discourse) about matters such as raining, being ready, being green structured according to a certain standard format. We have

seen how conceptual constraints help intuitive understanding, i.e. knowledge of the truth-conditions of an utterance in any arbitrary context. However, conceptual constraints could be at work even in our comprehension of utterances as performed in *particular* contexts. I shall explore this idea in the present section.

The idea I would like to defend is that, when sentences like "It's raining", "Jill is ready", "The leaves are green" or "Sam cut the cake" are at issue, it is reasonable to suppose that not only our intuitive comprehension, but also our comprehension of utterances *in particular* contexts relies on the non-context-specific capacities that underlie conceptual competence. If this is the case, then we have a model of how comprehension of these sentences works which envisions some *very basic constraints on pragmatic processes of enrichment*.

Suppose a speaker utters "It's raining", meaning that it's raining in Paris. Here is how the hearer could exploit her own conceptual (semantic) competence in order to work out the content of this utterance. First of all, the hearer needs to presuppose that the speaker is a competent user of at least the concept TO RAIN (if not of the predicate "to rain" as well). Secondly, the hearer needs to be aware that it is a conceptual constraint on every correct use of the notion of raining (and of the related predicate) that rain falls at some location: since the present utterance of "It's raining" is supposed to be a correct use of the concept TO RAIN (and supposedly, of the predicate "to rain" too), then the utterance must be true iff it's raining at some location x. At this point, the hearer knows that what the speaker has said is *that it's raining at some location x*. And, what is more, she has achieved this piece of knowledge only by applying the constraint to the effect that rain has to fall somewhere to the truth-conditions of an utterance of "It's raining". Of course, the story cannot end here. The hearer so far just knows that what the speaker is saying is true iff it's raining at some location x (salient to the speaker). Now, if that is possible, he needs to figure out *where* the rain is falling, i.e. he needs to provide a value for the x.

Here conceptual constraints cannot be of any help: the hearer *must* resort to speculation about the speakers' intentions. It is at this point that pragmatic-inferential reasoning has to become involved: at the level of fixing the reference to the particular place the speaker is talking about. What the hearer has achieved so far is a conceptual structure such as *It's raining in some x*; the next step is to provide a value for the location variable figuring in the conceptual structure. Here the hearer could engage in a pragmatic inference about reference to which location would make the speaker's contribution optimally relevant, or cooperative and so on. But note that engaging in pragmatic inference *at this point*, only for reference-fixing, is not particularly exciting. Fixing this element only entails establishing some background factor, which helps one to take only one of the many steps in the interpretive process—the other steps all being pre-determined and guided by one's mastery of the relevant conceptual (semantic) constraints. Enrichment processes appear as constrained by one's conceptual (semantic) competences, and pragmatic processes turn out as subordinated to the proper carrying out of these conceptually (semantically) constrained inferences.

We have seen an example of how conceptual (semantic) constraints may guide the reconstruction of the content of utterances of semantically under-determined sentences even in *particular* contexts. If things really work like this, what role is left for allegedly "free", pragmatically driven processes of enrichment? How much of *pure* pragmatics (as e.g. considerations of relevance, informativeness, perspicuousness, clarity, cooperation) do we put in our understanding of utterances of certain sentences, and how much do we rely on pre-determined constraints[22]? The considerations advanced

---

[22] "Pure" pragmatic processes seem to be involved in the cases of "shallow" semantic under-determinacy. Consider how Radical Contextualism conceives the way hearers reconstruct the content of an utterance of, for instance, "It will take time to get there": here the aim seems to be that of making the utterance's content as *relevant* and *informative* as possible—not of compensating come semantic defectiveness. The same applies for the way hearers may be thought to interpret utterances of "I've had breakfast", "Jack and Jill got married", "You are not going to die": in all these cases, the interpretive process is not guided by a demand for a semantically complete, evaluable content, but rather by a demand for a content which is most relevant, informative, perspicuous given the current aims of the conversation. Thus, the conceptual constraints view differs from Radical Contextualism in that it

in the previous sections suggest that the role of purely pragmatic inference may have to be reconsidered for the group of semantically under-determined sentences we have been investigating. In figuring out the truth-conditions of utterances of "It's raining", "Jill is ready", "The leaves are green", "Sam cut the cake" it seems reasonable to conceive that language users rely on pre-determined conceptual (semantic) constraints which in turn guide and pose limitations to pragmatic inferences concerning conversational cooperation, relevance and the intentions of conversants. For every use of "ready", "rain", "green", "cut", it is reasonable to countenance a related *conceptual (discoursive) template*, on which the hearer relies in construing the utterance's truth-conditions in the particular context in which the utterance is performed. Inferential, context-specific pragmatic processes may indeed play a role, however their role might be just that of "*filling in*" these pre-determined templates.

The resulting picture is by all means an alternative account of how utterance content is worked out by a hearer, where pragmatic-inferential processes play a much more modest role, and non-pragmatic, conceptual (semantic) factors do most of the job. An additional difference between Radical Contextualism and the present account becomes apparent if we switch for a moment to a conception of propositions as structured entities containing ordered *n*-uples of items of various sorts —e.g. individuals, properties. In the Radical Contextualist's framework, processes of enrichment introduce items in the proposition directly, with no mediation, only with the help of inferential mechanisms. In the present account, processes of enrichment don't introduce items in the proposition directly, because *a slot* (albeit a conceptual/semantic one) for these items must be pre-determined at a conceptual (semantic) level of representation. The location of the rain in the proposition expressed by an utterance of "It's raining" performed while talking about Paris is introduced in the proposition as the value of an un-filled position in the conceptual (semantic) structure associated with the utterance; the process of enrichment is therefore *mediated* by conceptual/semantic constraints. This may at first sight seem in sync with Sperber and Wilson's idea of "development of conceptual structure" and also with Recanati's idea of "expanding a sentence's syntactic structure" via variadic functions. Note, however, that what I am pointing at is something which in Sperber and Wilson's and Recanati's account goes almost unnoticed: namely, the fact that the "development" or "expansion" itself of the sentence's logical/conceptual form is *structured, or constrained*. Saying that conceptual structures are developed or expanded is not yet saying that these conceptual structures are developed *according to certain precise and pre-determined conceptual/semantic constraints*.

The advantages of adopting an account in terms of conceptual constraints are essentially three. Firstly, compositionality is rescued *at least at the level of the conceptual/semantic representation* associated with semantically under-determined sentences. Utterances of sentences like "It's raining" have been thought of as expressing contents which include elements that are not articulated by any component of the sentence (see Perry 1986), thus violating the principle of compositionality, which states that the content (in context) of a sentence is determined by the content (in context) of its components. We may, however, regard conceptual constraints as determining a further level of "conceptual representation" in which compositionality is not violated. The picture could be outlined as follows: uses of the semantically under-determined sentence "It's raining" trigger a conceptual representation whose form is *It is raining at some x,* where the locational position *x* may be assigned a value in context. Since the location that enters at the propositional level is the value of a variable which is present in the conceptual/semantic structure, compositionality is rescued to the extent that the content (in context) of the sentence is determined by the content (in context) of the components of the *conceptual/semantic representation* associated with an utterance of that sentence. The second advantage consists in the fact that an approach in terms of conceptual constraints accounts for utterance comprehension without portraying hearers as effecting any kind of pragmatic

---

envisages processes of contextual enrichment which are guided primarily by conceptual/semantic demands rather than by purely pragmatic demands for relevance, informativeness, etc.

inference, involving creative and explanation-oriented strategies, thus avoiding an *over-intellectualisation* of processes of utterance comprehension. It is not obvious that, when interpreting an utterance of "It's raining", hearers need to engage in any inferential reasoning as to which content would be more relevant, or about the communicative intentions of the speaker (see Recanati 2002 on this point). The processes hearers may go through could be much more "mechanical" than those envisaged by Truth-conditional Pragmatists. In the present account, following a conceptual constraint is indeed a "mechanical" matter which derives from one's conforming oneself to a pre-determined standard. What is regained in the present account is therefore the idea that comprehension of certain utterances is a rather effortless matter, which requires no particularly creative or sophisticated tasks from hearers apart from abiding by pre-established patterns of use.

The third advantage lies in the fact that the conceptual constraints view doesn't aspire to be a *syntactic* theory, i.e. a theory that says something about the real, perhaps hidden syntactic structure of some expressions in the language, like so-called "Indexicalist" accounts *à la* Stanley (2000) (which I briefly canvassed in the **Introduction** and in **section 1** of this chapter). The evidence coming from intuitive comprehension (i.e. knowledge of truth-conditions in arbitrary contexts) doesn't in any way license the postulation of syntactically realised, even though aphonic elements in the deep syntax of sentences. By contrast, it licenses a view in which the richer proposition is determined via the conceptual/semantic structures associated with the sentence's meaning. In other words, in the the conceptual constraints account processes of enrichment are not traceable to any element in the logical form of sentences, even though they are traceable to constraints posed by the correct use of concepts/predicates.

11. The Conceptual Constraints View and Formal Semantics

Is the conceptual constraints view compatible with a Davidson-style program of constructing a truth-conditional semantics for natural language? To adopt a picture in which utterance truth-conditions of semantically under-determined sentences depend on the speakers' conceptual/semantic competence doesn't imply abandoning a Davidsonian account of a theory of meaning. Truth-conditions can always be assigned compositionally to e.g. "Jill is ready" by a Tarski-style theory of meaning. What the present proposal adds to the level of a formal, compositional semantics is a level of purely *conceptual/semantic constraints on interpretation*. In other words, a truth-conditional, compositional and recursive semantics can be preserved, because it is only by means of such a system that we can maintain idea that sentences express contents on the basis of what they strictly speaking *mean*. However, formal semantics may deliver outputs which do not entirely fit the standard cognitive format through which thought and discourse about the relevant matters is organised. That is, the results of the compositional semantics of linguistic items may fall short of meeting certain requirements imposed by cognitive/discursive structures. This is where conceptual (semantic) constraints become helpful: their role is precisely that of adapting the deliverances of formal semantics to standard ways of organising thought, reasoning and discourse.

These considerations suggest the following picture. On the one hand, a formal semantic system matches an utterance of "It's raining" with its truth-conditions by deriving a T-sentence like "'It's raining' is true iff it's raining". However, familiarity with the concept TO RAIN and with the standard cognitive format associated with it allows the language-user to know that, since rain invariably happens to fall somewhere, uses of "It's raining" are true just in case it is raining *somewhere*. In general, one may envision a very "minimal" level of linguistic competence consisting in the capacity to derive *trivial* T-sentences like "'Snow is white' is true iff snow is white", in virtue of one's (tacit or dispositional[23]) knowledge of a formal semantic theory's axioms. This may not be

---

[23] There is a debate on what knowledge of a semantic theory consists in. Dummett's (1976) leading idea is that the ability underlying the use of language is tacit knowledge of propositions. The tacit character of language mastery has never been questioned, even though the idea that it consists in propositional attitudes has been criticised. For Evans (1981), tacit semantic knowledge is not a propositional attitude, but it is the possession of dispositions to issue judgments concerning truth, reference and satisfaction; for Davies (1987), tacit knowledge is possession of

sufficient in order to know how the world should be in order for utterances of that sentence to be true. What plays a role in enabling language-users to know what should be the case in order for *utterances* of (at least a certain group of) sentences to be true is what I shall call *conceptual/semantic competence*. In this sense, conceptual competence could be regarded as the capacity to derive *non-trivial* T-sentences such as "An utterance of 'Jill is ready' is true iff Jill is ready for some activity". This distinction seems to be acknowledged by Dummett, who distinguishes between a *theory of reference* and a *theory of sense*.

> The theory of sense specifies what is involved in attributing to a speaker a knowledge of the theory of reference. When the theory of reference takes the form of a theory of truth [e.g. a Tarski-style theory], this is necessary whenever an axiom of T-sentence assumes a trivial form, and therefore fails to display in what the speaker's implicit knowledge of it consists. (Dummett 1976: 127)

On the one hand, the theory of reference assigns a T-sentence to every well-formed sentence of the language, and could thereby be seen as describing the "minimal" level of linguistic competence; on the other hand, the theory of sense tells us what it is to *know* the truth-conditions of a sentence of the language, and could thereby be seen as describing what it is to know the truth-conditions of some sentences *also* in virtue of the exercise of one's conceptual/semantic competence. The two theories and the related competences do not seem to interfere with each other, so there seems to be no incompatibility between a project that pursues a description of the former and a project engaged in investigation of the latter, as long as their task is that of capturing different aspects of linguistic competence and, ultimately, two different levels of content.

12. Implementing Conceptual Constraints?
Throughout the chapter, I have assumed that the conceptual constraints view is intelligible as an autonomous proposal. It could however be objected that conceptual constraints are really *mysterious entities*, which neither belong to the linguistico-semantic realm, nor belong to the pragmatic realm. It would thus be an interesting issue to explore, *whether the conceptual constraints view should be reduced to any other theory* which makes clearer assumptions as to the representational level at which it operates. We have already ruled out, throughout the present chapter, that the application of conceptual constraints corresponds to any *purely* pragmatic process of the kind postulated by proponents of Truth-Conditional Pragmatics such as Carston and Sperber & Wilson. What remains to be explored is whether conceptual constraints can find any representation at either a semantic, or a syntactic, or even a sub-syntactic level. In this section, I shall simply gesture at three ways in which conceptual constraints could be represented, reserving extended discussion of these three alternatives for the next chapters.

The first possible account sees conceptual constraints as *hidden indexical elements*. As a matter of fact, this account would correspond to the view put forward by Stanley 2000 and Stanley & Szabó 2000, according to which contextual effects on content must find a representation in logical form: the account would therefore reduce conceptual constraints to *syntactic* entities. For example, the conceptual constraint to the effect that rain has to fall at some place or other shall be represented at the level of the logical form of the predicate "to rain" as an extra argument place. The real logical form of "to rain" thus becomes *rain(x, y),* where the locational variable y signals the presence of the constraint. The same goes for the predicate "being ready". The conceptual constraint to the effect that being ready is invariably being ready for some activity is represented at the level of logical form as an extra argument-place, so that the real logical form of "being ready" is *ready(x, y),* where the *y* variable ranging on purposes or activities is a mark of the conceptual constraint. Finally, the

---

explanatory states, whose causal structure is reflected by the derivational structure of semantic axioms and theorems (for an overview, see Miller 1997).

conceptual constraint that being green is being green in some respect is represented as an extra argument slot in the logical form of the predicate "being green" in such a way that the real logical form of the predicate becomes *green(x, y)*. I will be concerned with the hidden indexicals view in **Chapter 3**.

A second possible semantic account of conceptual constraints doesn't regard them as present at the level of logical form or syntax, but at the *lexical level*. The structure of lexical entries is to be thought about in terms of their *thematic structure*. A thematic structure is a general way of conceiving the terms of the relation introduced by a certain linguistic expression, according to very broad categories such as the agent/patient category. For example, the verb "kiss" expresses a relation that requires two different roles: that of the person who kisses and that of the person who is kissed. The kisser takes the thematic role AGENT, while the kissee takes the role THEME (or PATIENT). If the lexical entries of the expressions subject to conceptual constraints can be seen as having a thematic structure, whose thematic roles may correspond to the constraints at issue, then it is possible to carve up conceptual constraints as thematic roles. The representation of conceptual constraints would thus take place at a *subsyntactic level*, as opposed to the syntactic level, to which e.g. hidden indexicals belong. Take for example the verb "to rain": is it possible to turn the conceptual constraint to the effect that rain has to fall at some location into a thematic role? Here is how Taylor illustrates this possibility:

> Take the verb 'to rain' as an example. The view […] supposes that the verb 'to rain' has a lexically specified argument place which is theta-marked **THEME** and that this argument place takes places as values. This is a way of saying that the subatomic structure of the verb 'to rain' explicitly marks rainings as a kind of change that places undergo. Now from the point of view of sentence level syntax such lexically specified parameters are what I call subconstituents rather than constituents. Though subconstituents need not be expressed as sentence-level constituents, they make their presence felt by "demanding" to be assigned a contextually supplied value. (Taylor 2001: 53).

Once again the lexical proposal will be considered in **Chapter 3**, as a potential competitor of the hidden indexicals view.

The third possible account may regard conceptual constraints as affecting the *truth-value* of sentences, rather than their syntactic or subsyntactic structure. This account would correspond to a form of Relativism about the truth-value of sentences like "It's raining" or "Jill is ready" or "The leaves are green", along the lines of the solution proposed by e.g. Lasersohn (2005), Kölbel (2002), Egan (forthcoming). The Relativist strategy consists in positing that the circumstances of evaluation of a certain sentence do not consist simply of a possible world parameter *w*, but of a couple *<w, e>* including a possible world parameter *w* and an extra parameter *e*. The nature of this extra parameter may well depend on conceptual considerations: for example, when one is dealing with the semantics of personal taste predicates such as "tasty", one may want to see the truth-value of sentences like "This is tasty" as relative to a couple *<w, s>* composed by a possible world *w* and a taste-standard *s*. When one is concerned with the semantics of epistemic modals like "might" and "must", the truth-value of a sentence like "The treasure might be here" is going to be made relative to a couple *<w, e>* of a possible world *w* and an epistemic state *e*. In light of these facts, how then could a relativistic account of conceptual constraints look like? The account could go as follows: since the conceptual constraint that regulates correct uses of the concept "to rain" requires that every instance of a rain-event takes place at a location, then the truth of an utterance of "It's raining" is going to be made relative to a possible world *w* and a location *l*. The truth-conditions of an utterance of "It's raining" are thus going to be spelled out as being that an utterance of "It's raining" is true at *<w, l>* iff it's raining at *w, l*. Similarly, in virtue of the conceptual constraints that regulate

correct uses of the concept "being ready", the truth of an utterance of "Jill is ready" is going to be relative to a possible world w and a relevant activity a. And finally, since the conceptual constraints on use of the predicate "being green" dictate that every instance of greenness be an instance of greenness in some respect, then the truth of a sentence like "The leaves are green" will be relative to a couple consisting of a possible world w and a respect r. The related truth-conditional clauses would state that, for instance, an utterance of "Jill is ready" is true at $<w, a>$ iff Jill is ready at w, a; and that an utterance of "The leaves are green" is true at $<w, r>$ iff the leaves are green at w, r. The relative truth proposal will be examined in **Chapter 4**.

The previous ones were clearly just rough sketches of the potential theoretical accounts one might be willing to explore, the single positions mentioned being the object of a detailed evaluation in the chapters to follow. My aim for the present moment was just to give a flavour of how conceptual constraints may be implemented at some level of one's semantic theory—either at the syntactic level, or at the lexical level, or at the level of evaluation—where all these implementations try to escape the characterisation of conceptual constraints as operating at some arguably vaguely conceived "conceptual/semantic level". *It may turn out that none of these accounts is viable, after all*. As a matter of fact, in the end one may prefer an account in which conceptual constraints find *no* articulation at any level of representation—whether syntactic, subsyntactic or at the level of the index of the circumstances of evaluation. In the chapters to follows, I will try to argue that none of these ways of reducing conceptual constraints is really preferable to an account in which conceptual constraints operate at an autonomous level. A detailed outline of my proposal will be reserved for **Chapter 5**.

Conclusion

In this chapter I have been dealing mainly with Radical Contextualism (RC), i.e. the view that linguistically unconstrained pragmatic effects operate on a sentence's content-in-context.

- First, I have presented (RC)'s leading theses and I have illustrated how pragmatic processes can affect sentential meaning in context, via supplementation of the sentence's content with unarticulated constituents and *ad-hoc* concepts;

- Radical Contextualism employs a special notion of RC-comprehension, which amounts to knowledge of the truth-conditions of a particular utterance in a *particular* context. I have shown that Radical Contextualism predicts that, without the possibility to run pragmatic inferences aimed at content-enrichment, there shall be no RC-comprehension of utterances of semantically under-determined sentences;

- I have then introduced the notion of *intuitive comprehension*, which corresponds to knowledge of the truth-conditions of *any* utterance in *any arbitrary* context;

- Radical Contextualism has as its consequence that, if no context is available, there can be *no intuitive comprehension* of semantically under-determined sentences because, outside of a context, no pragmatic inference can be fruitfully run, and the hearer has to stay content with the sentence's meaning-in-context. Since this doesn't determine any truth-condition whatsoever for its utterances, there can be no comprehension of what any utterance of that sentence in any context expresses;

- Contrary to the Radical Contextualist's predictions, I stress the fact that, even outside of contexts, speakers *do* have an intuitive comprehension of utterances of semantically under-determined sentences. For example, speakers may know outside of a context that an utterance of "Jill is ready" is true in any arbitrary context iff Jill is ready *for some activity*; they may know outside of a context that an utterance of "It's raining" is true in any arbitrary context iff it's raining *in some place or other* and so on;

- In have proposed to regard the intuitive comprehension of utterances as depending on *conceptual competence*, i.e. familiarity with what counts as correct and incorrect uses of

a concept. Conceptual competence arises from a speaker's participating to a practice of concept use, which in turn leads subjects to acquire a standard cognitive format for thought on specific matters. Conceptual competence imposes *conceptual constraints* on thought and discourse. For competent language-users, conceptual constraints *de facto* coincide with semantic constraints. These constraints guide utterance comprehension of semantically under-determined sentences. For example, the conceptual/semantic constraint according to which rain necessarily falls at a location helps the hearer to understand an utterance of "It's raining" as being true iff it's raining *somewhere*.

- Not only do conceptual constraints help hearers to gain intuitive comprehension, i.e. comprehension of what any utterance of a sentence would say in any context. I argue that conceptual constraints also guide comprehension within a particular context. Thanks to conceptual constraints, the hearer *expands* the conceptual/semantic structure of the sentence (e.g. going from *<it's raining>* to *<it's raining in x>*) and fills in the newly opened semantic/conceptual gaps with contextual information (e.g. going from *<it's raining in x>* to *<it's raining in Paris>*).

- The conceptual constraints proposal is set apart from Radical Contextualism in that it envisages processes of enrichment which are *guided primarily by conceptual/semantic demands* and not by purely pragmatic demands of relevance, informativeness, etc.

- The conceptual constraints proposal is also compatible with a Davidson-style project of constructing a formal semantics for natural language. The proposal adds to the level of formal, compositional semantic interpretation a level of interpretation which proceeds from the language-user's conceptual/semantic competences coming from a standard cognitive/discoursive format.

- I have sketched some prospects concerning how to implement conceptual constraints into a formal semantic theory: the first prospective account makes use of the notion of *hidden indexical*; the second account takes advantage of the idea of *thematic structure*; the third account sees conceptual constraints as relativisations on *truth-value*. I shall be concerned with a critical assessment of these options in the following chapters.

Chapter 3

Indexical Contextualism and Conceptual Constraints

In the previous chapter, I introduced the problem of how comprehension of semantically under-determined sentences is possible. The problem, stated in **section 2.1** as the *[Comprehension Problem]*, concerns the general mechanism by means of which hearers manage to comprehend utterances of semantically under-determined sentences. In **Chapter 2**, I engaged in a criticism of Radical Contextualism: in my view, Radical Contextualism has a too strong notion of utterance comprehension, in that it envisages utterance comprehension as accomplished only in a particular context. Yet the theory fails to capture a notion of utterance comprehension in any arbitrary context, which I called Intuitive Comprehension. The intuitive comprehension of semantically under-determined expressions doesn't require any particular context to be accessible to the hearer. Language users usually understand utterances of "It's raining" and "The leaves are green" perfectly well, even when they are not acquainted with any context of utterance. The explanation I suggested for this is that hearers are guided by what I called *conceptual constraints* (which are, for competent language users, also *semantic* constraints). Conceptual constraints are ways of structuring thought about certain matters, such as episodes of rain, states of readiness. Employing conceptual constraints in the interpretation of utterances means adjusting the content of that utterance to a *standard format* which holds for thought and discourse with respect to that matter. I defended the idea that even processes of "enrichment" which take place within a specific context are guided by conceptual/semantic constraints, instead of being pragmatic through and through. That there are processes which are conceptually/semantically constrained marks a distinction with Radical Contextualism, for which enrichment processes are pragmatic in nature.

The claim that conceptual constraints play a role in the comprehension of utterances in specific contexts immediately elicits a question as to what the *status* of these conceptual constraints precisely is. As I just suggested, conceptual constraints are to be characterised as ways of cognitively organising information and hence, thought and discourse about the world. A claim like this however risks sounding too simplistic and commonsensical, or even mysterious.

In order to give a less mysterious status to what I have called conceptual/semantic constraints, one may prefer to talk about *syntactic* constraints. Appealing to syntactic constraints may seem to guarantee a more "compact" explanation of how utterance comprehension is possible in cases of under-determinacy, for syntactic constraints undoubtedly fall under the domain of linguistic competence; comprehension of utterances of semantically under-determined sentences could be explained only by appealing to the structure of the hearer's linguistic competence—rather than by appealing to the hearer's linguistic *plus* conceptual/semantic competence. More than one author in the last decade has put a strong emphasis on the alleged role of syntactic constraints in the comprehension of semantically under-determined sentences like "It's raining", "Jill is ready", "The leaves are green" etc. (see Stanley 2000, Stanley & Szabó 2000, King & Stanley 2005), and their reasons to do so had definitely to do with maintaining an explanation of so called "pragmatic effects" which could be related with linguistico-syntactic factors. In this chapter, I shall therefore explore the possibility that what I have introduced as conceptual constraints are really *syntactic* constraints on interpretation.

With respect to the nature of conceptual constraints, one could therefore go in the direction of an Indexicalist proposal, cashing out conceptual constraints as *hidden unsaturated elements in the*

logical form of expressions: for example, the conceptual constraint on locations in "to rain" is implemented as a hidden argument-place in the logical form of the predicate, so that "to rain" has the logical form *rain(t, l)*, where *t* is a time variable and *l* is a location variable. In this chapter, I will firstly present two versions of the position (**section 2**); subsequently, I will take into consideration the *evidence* that has been displayed in favour of it (**section 3**). My aim will be to show the weaknesses of the arguments that have been used to purportedly prove that an Indexicalist account is the only viable one, at least for certain classes of predicates (**sections 4-9**). I will then outline an alternative proposal which, if it is plausible, contributes to further undermining the case provided by Indexicalists (**sections 10-11**).

1. Overview

The Indexicalist approach gives rise to a precise semantic theory, called Indexical Contextualism (see MacFarlane 2009 and Recanati 2006 for use of this terminology). In **section 1** of **Chapter 2** I indicated Indexical Contextualism as the view that an expression's meaning is stable, but it determines a *function* from a context of utterance to a content. This is precisely what is expected to happen once we start to regard e.g. "to rain" as containing a hidden indexical. An Indexical Contextualist will say that "to rain" has a stable meaning, corresponding to a function from locations to properties. For every context in which "to rain" is uttered, the function maps the contextually salient location *l* to the property of *raining-in-l*. Thus if "It's raining" is uttered in Paris, while talking about the weather in Paris, the function associated with "to rain" maps the contextually salient location, Paris, to the property of *raining-in-Paris*; the content of the utterance will thus be *that it's raining in Paris*. Similarly, if "green" is regarded as covertly indexical, containing a hidden variable for ways or "respects" of being green, then for every occurrence of "green" in a context, the function associated with the predicate will map a contextually salient respect *r* of being green onto a property of *being-green-in-respect-r*. So for example, when Pia utters "The leaves are green" in a context in which being green on the outer surface is salient, the function associated with "green" will map the salient "on-the-outer-surface" respect onto the property of *being-green-on-the-outer-surface*, thus making the content of Pia's utterance *that the leaves are green on the outer surface*.

Indexical Contextualism has been chosen as an option to represent the semantic profile of a wide array of expressions, for many different reasons. Thus for instance, the semantics of a term like "know" has been modelled along the lines of indexicals, (see Cohen 2000, DeRose 1995), in order to explain skeptical paradoxes. So, for instance, an assertion of "S knows that *p*" in context *c* expresses, for these theorists, the content *<that S knows that p by the epistemic standards of c>*. Moving to a different context *c'* may have the effect of changing the operative epistemic standards —by raising them, for example. This entails that the content of an assertion of the form "*S* knows that *p*" becomes *<that S knows that p by the epistemic standards of c'>*. Since considering skeptical options (such as the brains in a vat hypothesis) is taken to raise the epistemic standards, common assertions like "I know I have hands" become false at the context in which the skeptical scenario is being considered. This, however, doesn't prevent that, in non-skeptical contexts, claims of knowledge are still true and felicitously assertible.

Moral terms like "good", "ought" (see Harman 1978) as well as predicates of personal taste like "delicious" or "disgusting" (see Glanzberg 2007, Cappelen and Hawthorne 2009, De Sa 2008, Sundell 2011) have been treated like indexicals too, in order to account for the fact that the truth of moral or taste judgments depends on the inclinations of the speaker (or of a group which is salient to the speaker). So for instance, the content expressed by a statement of "Slavery is morally wrong" in context *c* is *<that slavery is morally wrong by the moral standard of c>*. The content expressed by an utterance of "Apple pie is delicious" in context *c'* is *<that apple pie is delicious by the gustatory standards of c'>*. Changes of context can determine changes of content and—obviously— also of truth-value.

Quantifiers like "all", "some", "the", as well as gradable adjectives (see Stanley 2000), nouns and noun phrases (Stanley and Szabó 2000) and adjectives in general (Szabó 2001, Rothschild and Segal 2009) have been given an Indexical Contextualist semantics in order to account for the apparent context-sensitivity of their use. So for instance, an utterance of "Every one was accepted to the conference" in context c expresses the content *<that every one in the salient domain of c was accepted to the conference>*. An utterance of "Every student failed" in context c' expresses the content *<that every student-in-the-c'-domain failed>*. An utterance of "Naomi is rich" in c'' expresses the content *<that Naomi is rich by the standards of c''>*. An utterance of "The leaves are green" in context c''' expresses the content *<that the leaves are green in the respect salient in c'''>*. What would an Indexical Contextualist option offer to the proponent of conceptual constraints, as opposed to Radical Contextualism? It may be observed that Indexical Contextualism doesn't give rise to the wrong prediction that—as I argued in **Chapter 2—**followed from Radical Contextualism, according to which we do not have an Intuitive Comprehension of semantically under-determined sentences out of a context. For if a sentence like "Jill is ready", "The leaves are green" or "Olivia is tall" is indexical, then that means that the involved expressions have a standing linguistic meaning, which is identifiable whether the expression is interpreted "in" or "out" of a context. Intuitively understanding these sentences will require only the exercise of the speaker's linguistic competence. Surely, understanding an utterance of such sentences as performed within *a particular context* will require more than mere intuitive understanding: the hearer shall have to look at the context of utterance, searching for what, in context, would satisfy the indexical character of the expression. But after all, this is compatible with the expressions at issue being indexical. Language users are not able to understand what, in a particular context, utterances of "I am the king of the world" or of "It's freezing in here" express, if they have no idea of what the referents of "I" and "here" are. Yet they may intuitively understand the same utterances, independently of the context, precisely insofar as they master those expressions linguistically. So Indexical Contextualism, unlike Radical Contextualism, doesn't predict that a token-sentence of "It's raining", heard out of its context of utterance, won't be *intuitively* understood by a competent language user, even though it maintains that comprehension of utterances within particular contexts require more than intuitive comprehension. It seems that comparatively with Radical Contextualism, Indexical Contextualism may offer an account which is more in line with our pre-theoretic views about utterance-understanding. But is Indexical Contextualism the right theory to endorse for the proponent of conceptual constraints? First of all, we need to explore in closer detail how Indexical Contextualism is motivated, how it is supposed to work and what (syntactic) evidence we have to endorse it. That's what I will be concerned with in the following few sections.

2. Two Ways of Rescuing Compositionality
We saw in the previous chapters that several authors argued that meaning (in context) under-determines truth-conditions; that is, that what a certain sentence conventionally says (modulo ambiguity resolution and reference-assignment to indexicals and demonstratives) is insufficient to specify the truth-conditions of any utterance of it, in a way which seems to have nothing to do with either ambiguity, vagueness or context-dependence (indexicality). This undermines the idea that semantics is concerned with truth-conditions: if the meaning of a sentence (in context) doesn't express its utterances' truth-conditions, then this means that semantics delivers non-truth-conditional objects, which Bach (1994a) has called "propositional radicals".
But these arguments do not only threaten the truth-conditional nature of semantics. They threaten also its *compositional* character. A semantics for utterances of sentences in a language L is compositional iff it has it that the content and truth-conditions of an utterance u of a sentence s are determined by the contents (in context, where appropriate) of $s's$ components and their syntactic arrangement. Here the notion of "determination" is crucial: it doesn't simply imply that the content and truth-conditions of an utterance depend on the content-in-context of its components, for this

would allow for the utterance's content to depend also on *other* factors. Determination is a stricter relation: it is such that the content and truth-conditions of an utterance must depend *exclusively* on the content-in-context of its components. Only compositionality accompanied by this strict notion of determination can serve the two main purposes that it has usually been taken to serve: to explain how speakers learn language (*learnability*) as well as how they understand sentences that they never heard before (*novelty*) (see Pagin&Westerstahl, 2010).

Semantic under-determinacy entails the failure of compositionality for utterance content, since if an utterance's truth-conditional content is not *exclusively* determined by the content (in context, where appropriate) of its components and their syntactic arrangement, this means that such truth-conditional content will have to be determined also by other elements—perhaps mechanisms of pragmatics completion or enrichment.

More than one author has come to question the principle of compositionality on these grounds, if not for the whole language, at least for a class of expressions: if these authors are right, language—or at least substantial fragments of it—is not compositional. For example, Lahav (1989) contends that adjectives such as "red", "sharp", "good" do not behave compositionally. Lahav draws attention on the fact that the semantic contribution, i.e. applicability condition, of adjectives is not determined once and for all for every occurrence of the words, but it is rather dependent on the noun with which they are combined. For example, when "red" is part of the compound "red pen", its semantic contribution seems to differ from its semantic contribution within the compound "red bird". This is because the conditions in which "red" applies are different when what is at issue is pens rather than birds. For a pen to be red (in normal circumstances), it doesn't need to be red all over: it's enough that it contains a red ink, such that it allows its user to trace red signs on a piece of paper. This conditions obviously differ from the conditions at which a bird is red, i.e. that the bird's feathers be red (even though not its beak and its legs). These considerations lead Lahav to the conclusion that the semantic functioning of adjectives is much more complicated than that described by a compositional semantics, because it has to make room for a large degree of contextual variation in their applicability conditions. Since any way of complicating their semantics would either push the problem one step further or postulate implausible semantic notions (endless ambiguity, "primitive" redness), this compels Lahav to abandon the idea that adjectives are ever compositional.

> Compositional adjectives, even if they are such, can no longer be accounted for in a straightforward way, as simple, isolable, semantic building blocks […]. […] [A]djectives should be construed as having a non-compositional semantics. (Lahav 1989: 278).

In response to this attack, more than one solution has been proposed, in order to preserve the idea that the semantics of adjectives *is* compositional. Some theorists have recently proposed that colour predicates like "green" and "red" are simply ambiguous (see Kennedy and McNally 2010). Others have opted for an "indexical way out": for these theorists, rescuing the compositionality idea requires accommodating some parameter for the contextual variation of the expression's content into the syntactic structure of the expression itself. In this section, I will focus on two "Indexicalist" proposals, dealing specifically with adjectives.

A first proposal is put forth by Szabó (2001). Szabó first considers Travis' "green leaves" example. In Travis' example (see **Chapter 1**), two utterances of the same sentence ("The leaves are green"), performed in the same worldly circumstances, have two different truth-values, but none of the sentence's constituents appears as obviously context-dependent. Travis then speculates that, even though no context-sensitivity is exhibited by the sentence, nevertheless the sentences express two different contents in different contexts. According to Szabó, this conclusion undermines what he calls the "Context Thesis", according to which

> The content of an expression depends on context only insofar as the contents of its constituents do (Szabó 2001: 122).

Since the content of "The leaves are green" changes across the two contexts, even though none of its constituents' content does, the Context Thesis is refuted. Szabó goes on and says that, since the Context Thesis derives from the Principle of Compositionality, then Travis' argument, by providing a counterexample to the former, counts as a refutation of the latter.

Szabó's way of restoring the Context Thesis and, consequently, an orthodox compositional semantics for utterances of expressions like "The leaves are green" consists in arguing that "green" is an *incomplete predicate*. The word "green" doesn't correspond to the monadic predicate GREEN(X), but to a two-place relation GREEN(C,P), in which C stands for a comparison class and P stands for a part in which the object is green. Szabó then suggests that

> the context-dependency that appears in Travis's example is of a relatively easily characterizable kind: it is a matter of different contextually specified values for the variable P. (Szabó 2001: 138)

That is, the content (and truth value) of the first utterance of "The leaves are green" is determined by fixing the value of the P-variable as something like *<on the external surface>*, while the content (and truth-value) of the second utterance is determined by fixing the value of P as being something like *<under the paint>*. By means of this formal modification of the syntactic structure of the predicate "green", one can explain the fact that the sentence "The leaves are green" has different contents and truth-values in different occasions of utterance by saying that the word "green" is context-dependent in the way of a (broadly construed) indexical. The context-sensitivity of the sentence is explained by the context-sensitivity of a constituent of the sentence (the word "green"), in such a way that the Context Thesis is upheld and a standard compositional approach for utterances of "The leaves are green" is thereby rescued.

Szabó's approach has been criticised recently by Rothschild and Segal (2009). These authors suggest an approach in which "red" (or "green") is not associated with a couple of variables in logical form, but it is simply indexed with respect to the context of utterance. As the authors claim:

> our semantics differs from Szabó's in that rather than positing indexical variables in the syntax introduced along with the predicate 'red' we treat 'red' itself as an indexical. (482)

Their proposal thus manages to

> posit a relatively simple syntax, with only a little more structure than meets the eye. In this respect our proposal is more economical than Szabó's, who posits two hidden variables. Our proposal is as simple as a theory that posits just a single hidden variable referring to the context of utterance and serving as a catch-all covering a very wide range of contextual variations. (474)

According to their account, out of a context, an adjective like "red" has just a generic intension red_g, corresponding to "red in some way or other", while it acquires a specific intension in each particular context. So, in the context c_j where "The pen is red" is uttered, the contribution of the word "red" to the content of the utterance is a more specific intension red_j; while in a context c_h where "The parrot is red" is uttered, the contribution of "red" will be a c_h-specific intension red_h. The proposal is, by explicit admission of the authors, the formal-semantics version of the Radical Contextualist's version of "ad hoc" or "pragmatically enriched" meanings:

It is worth noting that our analysis is similar to another, more pragmatic treatment, of 'red'. On this treatment, 'red' has a constant, context-independent semantic value, which value gets enriched in a given context of utterance (see for instance, Recanati[24] [...]). The newly enriched value then combines compositionally with the rest of the elements in the sentence in the usual way. This theory, unlike ours, gives a genuine role to pragmatic enrichment (rather than indexical resolution) in the determination of meaning. On the other hand, the pragmatic enrichment plays the same formal role as the indexical resolution does in our […] theory, so there is little difference in the basic structure of the accounts. (473)

It is not entirely obvious that Rothschild and Segal's proposal marks an improvement over Szabó's theory. The reason resides in the rather loose constraints posed on context by the indexical element posited in their account. Unlike Szabó's indexical element, whose task is that of selecting a precise contextual feature such as the part P and the comparison class C, Rothschild ans Segal's "catch-all" indexical element seems *too generic*: what aspect of the context is it supposed to focus on? There seems to be no answer available except for: any aspect that is salient to the speaker who performs the utterance. However, this way of conceiving of an indexical element seems to overlook the fact that *indexicality doesn't perfectly coincide with contextual saliency, but it rather puts constraints on contextual saliency*. Even if we consider indexicals whose reference is fixed by means of saliency criteria such as "this" and "that", we may observe that each of these expressions constrains saliency considerations in its own way. "This" constrains saliency considerations by restricting individuation of the salient object with respect to the set of objects that are *close enough* (both physically and metaphorically, I suspect) to the speaker, while "that" directs saliency consideration towards objects that are *far* from the speaker (again, both physically and metaphorically). In other words, the semantics of these expressions guides—even though in a loose way—the saliency-based individuation of their referent. This doesn't happen with the "catch-all" parameter postulated by Rothschild and Segal, which looks like a place-holder for whatever aspect of the context may prove relevant. The following question then arises: is the economy in terms of postulated parameters a genuine gain, if contrasted with the loss of specificity of the single index proposed by the authors? These considerations lead me to the conclusion that the two proposals are *prima facie* on the same level as far as the costs and benefits that they supply. I shall therefore consider them as substantially equivalent versions of the same formal template, which is based on the idea that the semantics of certain expressions is best cashed out in terms of the presence, in their logical form, of elements that syntactically pin-down the context-dependence of their content. The motivation behind both proposals is, invariably, to rescue the compositionality of utterances of a certain group of expressions, whose compositional nature has been put into question. In the next section, I will delve more into the issue of *what evidence* there exists for opting for an Indexicalist solution.

3. The Evidence
So far I have been concerned with the underlying motivation in favour of an "Indexicalist way out"

---

[24] Curiously enough, Recanati himself, who is a leading proponent of pragmatic enrichment, recently made a proposal aimed at reconciling pragmatic enrichment and compositionality whose outline is very much in the same vein as Rothschild and Segal's proposal, which he terms as a "Modulation Account" (see Recanati 2010, but also Pagin and Pelletier 2007 for a similar account). As he holds: "Even though free pragmatic processes are allowed to enter into the determination of truth-conditional content, still [...], they come into the picture as part of the compositional machinery. Semantic interpretation remains grammar-driven even if, in the course of semantic interpretation, pragmatics is appealed to not only to assign contextual values to indexicals and free variables but also to freely modulate the senses of the constituents in a top-down manner. Semantic interpretation is still a matter of determining the sense of the whole as a function of the (possibly modulated) senses of the parts and the way they are put together." (22)

of the problem of semantic under-determinacy and pragmatic intrusion: rescuing the compositionality of utterance content. I have also sketched two (at bottom equivalent) proposals aimed at making the semantics of adjectives like "red" similar to indexicals, in order to preserve the idea that they are compositional. But what is the *evidence* that allows such a move? *Are there any evidential and argumentative grounds for embracing an Indexicalist position*? In this respect, the most accurate survey of the evidence for Indexical Contextualism has been put forward in Jason Stanley's "Context and Logical Form". In this paper, Stanley intends to show that "all truth-conditional effects of extra-linguistic context can be traced to logical form" (Stanley 2000: 391). His critical aim is the unarticulated constituents account inspired by Perry (1986), which was covered in **section 4** of **Chapter 2**, according to which context provides content with elements that have no correlate in the sentence's syntax (or logical form).

Stanley can be regarded as engaged in an attempt to both to save compositionality and to rescue the idea that semantics is still concerned with truth-conditions, rather than with non-truth-conditional entities such as "propositional radicals", that pragmatic mechanisms merely flesh-out via the provision of unarticulated constituents. In order to do this, and since "claims of unobvious context-dependence must […] be accompanied by arguments for the existence of a corresponding formative in the logical form of the relevant constructions" (401), Stanley tries to bring independent syntactic evidence for the view that sentences like "It's raining", "Olivia is tall" etc. are really indexical, and so they express genuine truth-conditional contents.

Stanley's first step is that of acknowledging the fact that sentences allegedly containing unarticulated constituents are subject to *bound readings,* in which "the interpretation of the alleged unarticulated constituent is controlled by an operator in the sentence" (411). Bound readings are purely semantic phenomena; nonetheless, Stanley endorses what he calls the *Binding Assumption* to the effect that, whenever there is a semantic binding, then there is also a syntactic binding: that is, "semantic binding and syntactic binding coincide" (412). This leads the way to Stanley's conclusion to the effect that, whenever there is a bound reading, then the allegedly *un*articulated constituent must be articulated, in virtue of the fact that binding as a semantic phenomenon must co-occur with binding as a syntactic relation, where syntactic items are articulated by definition.

Let us report Stanley's examples. His initial focus is on "It's raining". According to the unarticulated constituents analysis, a simple occurrence of "It's raining" has the logical form "It is raining *(t)*". However, this logical form doesn't correctly represent the content of an utterance of "It's raining", in that what utterances of this sentence usually express is *that it is raining at a certain time* t *at a certain location* l. Here, theorists such as Perry contend that the location figures in the expressed content as an unarticulated constituent. An utterance of "It's raining" will thus be true iff it is raining at a certain time *t* in a certain location *l*, where *l* is an unarticulated constituent of the (structured) proposition expressed by the utterance. Stanley thinks that this analysis is incorrect. For consider the following sentence

(1) Every time John lights a cigarette, it rains.

The proposition expressed by this sentence is very plausibly *that every time John lights a cigarette, it rains in the location* l *where John lights the cigarette*. The proposition exhibits a semantic binding in the sense that the quantified phrase "Every time John lights a cigarette" seems to introduce a restriction upon the space-time locations at which it rains—viz. the restriction that these be the space-time locations in which John is lighting a cigarette. Since there is a semantic binding, then there must be a syntactic binding, i.e. the operator introduced by "every time …" must bind a variable *l* ranging over locations (besides obviously binding a variable *t* ranging over times). But if the location *l* is an unarticulated constituent, then it cannot be bound, since (as it follows from the previously endorsed *Binding Assumption*) binding operators only bind elements which are articulated in the syntax.

The result would be that an analysis in terms of unarticulated constituents would only predict an *unbound* reading of (1), namely the reading *that every time John lights a cigarette, it rains in some contextually salient location l.* Since the the bound reading is also an available reading for (1), but an unarticulated constituent analysis doesn't capture it, Stanley concludes that an account in terms of variables in logical form is preferable, since it maintains the possibility for both the bound and the unbound reading. Stanley thus illustrates the advantage:

> The problem with the unarticulated constituent analysis is that it only predicts one of the two available readings for [(1)]. In contrast, an account involving the postulation of a location variable predicts both readings. For variables can either be bound or free. An account involving variables, therefore, predicts there to be two readings of [(1)], one in which the value of the relevant variable is supplied by context [...] and one in which it is bound [...]. (417)

To sum up, Stanley argues that an analysis that postulates a variable *l* for locations in the logical form of (1) allows to capture both the readings that are available for (1): (i) the reading *that every time John lights a cigarette, it rains in the location l where John lights the cigarette*, in which the "every-time" operator binds the variable *l*; (ii) the reading *that every time John lights a cigarette, it rains in a salient location l*, in which the variable *l* is left free. This is a reason, according to Stanley, to adopt a view according to which sentences like "It's raining" contain a hidden variable in logical form.

The argument could be replicated also for other expressions, for which an account in terms of unarticulated constituents has been (or could be) provided. Besides "It's raining", Stanley focusses also on comparative adjectives ("tall" or "small"), quantifiers ("all", "some"); and relational terms such as "home", "local", "enemy"; in principle, his arguments can be extended to all those cases in which an unarticulated constituent theorist could propose his treatment (hence potentially also to terms like "green" or "ready"). Let us see some of the further examples Stanley presents.

The next class of expressions that Stanley surveys in order to defend his Indexicalist approach is that of *comparative adjectives*. Sentences such as "Snowball is small" can change their truth value according to what class of comparison is salient. Suppose that Snowball is an over-size cat. If the current class of comparison is elephants, then it is true that Snowball, being a cat, is small; however, if the class of comparison were other cats, it is false that Snowball, as an over-sized exemplar, is small. Theorists such as Bach (1994a), or Heim and Kratzer (1998) explain this shift in truth-value by saying that the content expressed by each different utterance of "Snowball is small" is different, in that it contains two different unarticulated constituents: in the first case, it contains the set of elephants as the unarticulated comparison class; in the second case, it contains the set of cats. If this explanation is right, then the semantic rule associated with "small" would have to be such that the denotation of "small" is the class of all things that are small with respect to the *contextually salient comparison class s*. Stanley contends that this kind of analysis can't be right. Take the following sentence:

(2) Most species have members that are small.

There are two readings that this sentence can take. As you may notice, the first reading is a bound reading, while the second is a free reading:

(3) Most species *p* have members that are small with respect to *p*;
(4) Most species *p* have members that are small with respect to a contextually salient comparison class *s*.

As a matter of fact, the first reading is not predicted by the unarticulated constituents theorists. That is because (by the *Binding Assumption*) bound readings require articulated variables to be bound by quantifiers; *un*articulated constituents clearly cannot be bound. So there is no chance that, if the comparison class is unarticulated, a comparison class-variable s is bound by the quantifier "Every species". However, since the first reading should be acknowledged as available just like the second, and since the unarticulated constituents account can't provide this reading, it follows that the account is defective. An account that regards "small" as containing a hidden comparison class-variable in logical form, such that it can be bound or left free, is preferable, because it predicts both readings.

The same argument can be repeated for *quantifiers*. It has been argued that utterances of quantified sentences often express propositions in which the quantification is restricted to a salient domain. This may be explained, for example, by saying that context provides a property that restricts the quantification. Thus:

(5) Every bottle is in the fridge

may express the proposition *that every bottle in the house is in the fridge*, where the property of being in the house restricts the domain of quantification of the quantified phrase "Every bottle". Some might think that quantifier restrictors are unarticulated constituents, because there need be no variable in logical form whose value is the restricting property. However, if this were true, bound readings could again not be predicted. For example, the following sentence

(6) In most of John's classes, he fails exactly three students

Sentence (6) has two available readings, the first bound, the second free:

(7) In most of John's classes x, he fails exactly three students in x;
(8) In most of John's classes x, he fails exactly three students in a contextually salient domain d.

The first bound reading is undoubtedly more plausible than the second unbound one, which barely makes sense. However, this is exactly the reading that is left out by the predictions of the unarticulated constituents view: for if the domain of quantification is not represented by a variable at the level of logical form, it cannot be bound by a quantifier. In light of these results, Stanley concludes again that the Indexicalist account is preferable, in that it predicts the desired bound reading. The same considerations have been provided by Stanley concerning relational terms such as "home" or "enemy", while Stanley and Szabó (2000) argue for an Indexicalist semantics concerning the logical form of nouns and noun phrases.

To summarise, in this section I have presented Stanley's main argument in favour of an Indexicalist approach, which has been called by some "The Binding Argument". In the following sections, I shall survey the main objections with which the argument has been attacked.

4. Objection 1: Inconclusiveness

The structure of Stanley's argument is the following: first, consider some complex constructions such as "Everywhere I go, it rains"; second, make sure that these complex constructions exhibit a bound reading, i.e. a reading in which the value of a location variable is controlled by another element, like a quantifier over locations. The *Binding Assumption* tells us that, since there is semantic binding, then there must be syntactic binding, for no binding operator binds an unarticulated element. So, the conclusion goes, there must be an articulated variable in the logical form of "Everywhere I go, it rains". So far so good. Stanley though, takes a further step, claiming that there must be a variable in the logical form of the *un-embedded* sentence "It's raining" too (see

Stanley 2000: 416).

The conclusion that *un-embedded* occurrences of the predicate "to rain" contain a location variable is, however, unwarranted, at least if we follow the explicit structure of Stanley's argument. In order to establish it, Stanley would need an extra premise. This premise presumably remains implicit in the way he runs his argument. Let us see how this extra premise, if made explicit, would interact with the other premises in the argument. The "missing" premise is in boldface in the following reconstruction:

1. In "Wherever I go, it rains" there is a semantic binding;
2. In order to have semantic binding, we have to postulate that there is a hidden variable in the embedded sentence "... it rains" (due to the *Binding Assumption*);
3. **If there is a bound variable in "... it rains", it means that this variable is in the predicate's logical form even when the expression occurs un-embedded;**
4. There is a bound variable in the embedded "... it rains";
5. So there is a hidden locational variable in the un-embedded "It's raining"

Premise (3) states that if there is (reason to postulate) a location variable in the embedded occurrence of "… it rains", then there must be a variable also in un-embedded occurrences of "It's raining". But is this necessary? What is the motivation for this assumption? A rather obvious motivation would be *uniformity*, i.e. the idea according to which a certain theoretical posit should hold for all *relevantly similar* cases, and not only for a restricted set of them. Uniformity is clearly a virtue of a syntactic theory, in that it protects the theory from the risk of being *ad hoc*. If, therefore, a variable in the logical form of "...it rains" has to be postulated when the expression is embedded, then by uniformity one should suppose that a variable has always been there, even though there was no operator to bind it, for all occurrences of "to rain", whether embedded or not, should count as relevantly similar. If that weren't so, the postulation of the variable in the logical form of the un-embedded clause "It's raining" may be regarded as an *ad hoc* move. But how motivated is the claim that an embedded and an un-embedded occurrence of "to rain" are relevantly similar?

There is an obvious respect in which the two occurrences are relevantly similar: they are both occurrences of the predicate "to rain" (disregard for the moment differences in tense). However, is this consideration strong enough? I argue that it may be overridden by considerations concerning *the linguistic context* in which the two occurrences take place. These considerations may undermine the thought that the two occurrences are relevantly similar, and thus block the extension of the hidden variable account to un-embedded occurrences of "to rain".

A first consideration may be to the effect that a construction like "Everywhere I go, it rains" exhibits a bound reading, but only because the explicit quantification over locations "Everywhere" causes the logical form of the embedded clause "... it rains" to be increased by one argument-place. This argument-place is occupied by a variable that gets bound by the quantifier itself. As Martì observes (but see also Recanati (2002a), Carston (2002) , Bach (2000)), it may be the case that, in "Everywhere I go, it rains"

> the quantifier takes *it rains* as its argument and gives back 'For every *l*. . .it rains *in l*' (i.e., the quantifier introduces the place of rain, fills this role with a variable, and binds it). (Martì 2006: 161)

This would explain why a variable is to be posited in the logical form of the complex sentence "Everywhere I go, it rains"; however, nothing would follow from this about the logical form of the *simple* sentence "It's raining". Since no operator is present in the simple sentence, whose role may be that of introducing (and at the same time binding) the location variable, the case would not be relevantly similar to the previous one. The move of extending the hidden variable- account would

not therefore be grounded by considerations of uniformity, on the account that an occurrence of "to rain" as embedded by an operator like "everywhere" and an occurrence of "to rain" as un-embedded do not count as relevantly similar in the present case.

A second consideration, even more radical than the earlier one, consists in noting that, even in the complex sentence "Everywhere I go, it rains", no binding occurs. This is because the phrase "Every where I go" is not a variable-introducing operator, but a simple "modifier" or "adjunct", which specifies the conditions at which the rain event described by the clause "... it rains" occurs. Since modifiers do not alter the adicity of the predicates they modify, this means that there is no good syntactic reason to postulate a hidden variable in the logical form of "Everywhere I go, it rains". But then, *a fortiori*, there is no good syntactic reason to postulate a hidden variable in the logical form of the simple "It's raining" (see Collins (2007)).

The considerations raised above are of significance to the extent that they cast doubt on the correctness of the passage, which is rather smooth in Stanley's reflections, from postulation of a hidden variable for the embedded "... it rains" to postulation of a hidden variable in the *un-embedded* "It's raining". The passage is licensed by the "implicit" premise **[3]**, which I have emphasised in the reconstruction of Stanley's argument. As a matter of fact, premise **[3]** is unwarranted, for the uniformity considerations that apparently underlie it can be undermined by considerations regarding linguistic context, like the ones I have exposed. Ultimately, then, Stanley's argument is inconclusive. The claim that a simple, un-embedded occurrence of "to rain" has a hidden argument in logical form, which lies at the heart of his strategy of "tracing all pragmatic effects to logical form", doesn't cogently follow from Stanley's argument, because it relies on an unwarranted (implicit) premise.

5. Objection 2 – Overgeneration

Overgeneration arguments play with the fact that bound readings such as that exhibited in (1) are very frequent in language interpretation. If this is so, the argument goes, then this should lead the theorist to postulate hidden indexicality in each of these cases. This would entail generating further hidden variables, which would make the expressions involved context-sensitive in a way which, however, goes counter to our intuitions. Generation of hidden positions is therefore *over*-generation, to the extent that the further variables appear to be not required at the intuitive level.

Arguments from overgeneration have been advanced by more than one author in the literature (see Cappelen and Lepore 2002, 2005, Recanati 2002a, Breheney 2003, Hall 2008), but it is important to distinguish between a "good" and "bad" version of the argument. Let us review both typologies.

Good arguments from overgeneration start with a *genuine* bound reading and conclude that, if the semantics is correct, we have to postulate more hidden syntax. This entails generating new hidden positions, thus acknowledging a dimension of context-sensitivity which runs against our intuitions. An example may be the following:

(9) Everywhere she goes, Mary dances;

In (9), it is plausible to suppose that—as far as the semantics goes—the embedded clause contains a variable *x* which is bound by the quantifier "everywhere", the resulting interpretation being <*that for every place x where she goes, Mary dances in x*>: indeed this seems to be required in order to make sense of the truth-conditions of the sentence. By the Binding Assumption, this would also require postulating a hidden variable in the *deep syntax* of the embedded clause "Mary dances". This means recognising that the logical form of the verb "to dance" has an extra argument place for locations, its structure being *dance(x, y)* rather than simply *dance(x)*. However, it doesn't seem that the predicate "to dance" is location-sensitive. The Indexicalist owes an explanation of why we have to admit a dimension of context dependency which goes against our semantic intuitions.

The argument may be regarded as a good one insofar as (9) involves a genuine example of semantic

binding. This can be tested in various ways: a first test consists in assigning different values to the bound variable, checking if the truth-conditions of the embedded clause change accordingly: it is quite easy to ascertain that it is so, for the truth conditions of the embedded phrase do change across "If Mary goes to Paris, she dances in Paris" and "If Mary goes to London, she dances in London" (see Cohen and Rickless 2007). A second test, proposed by Sennett (2008), has it that (9) is an instance of semantic binding just in case it is possible to felicitously deny the binding interpretation. Indeed, it seems felicitous to reply to an utterance of (9) with the following: "No she doesn't—she dances just in *some* places". Another test proposed by Pagin (2005) consists in evaluating the semantic acceptability of bound readings which contain an explicitly bound element, such as a pronoun ("it", "there")), in place of an element which is only implicitly bound (a covert variable x). In the case of (9), it seems entirely semantically acceptable to state "Everywhere she goes, Mary dances *there*".

"Bad" arguments from overgeneration start with something which is *not obviously* a bound reading and reach the same conclusion as the "good" arguments. However, since the reading which they exploit is not a genuine binding (or it is dubious that it is), then the whole argument is unsound. Examples of "bad" versions of the binding argument may be the following:

(10) In some way or other, the policeman stopped the car;
(11) Wherever I go, 2+2 = 4;

Recanati (2002a) brings (10) as an example of bound reading. However, Pagin (2005) questions the fact that a genuine semantic binding is in place. The reason is that, once the alleged implicit bound variable is turned into an explicit pronoun (of the same kind as "it", "there", "that"), as in "In some way or other, the policeman stopped the car *in that way*" the sentence is semantically odd, in the sense that the semantic binding is not immediately evident, nor it is required in order to make sense of the sentence's truth-conditions. Turning to (11), Cappelen and Lepore (2002, 2005) hold that applying the binding argument here results in postulating a hidden variable in the logical form of the sentence "2+2=4", which sounds utterly counterintuitive. Several authors, however, have questioned their assumption that (11) is an instance of bound reading. Pagin (2005) argues that, since "Wherever I go, 2+2 = 4" and the isolated sentence "2+2 = 4" cannot differ in truth-value, there is no semantic clue from which one can derive a location-dependence (i.e. a semantic binding) in (11). Similarly, Cohen and Rickless (2007) hold that it is implausible to suppose that the logical form of (11) is as follows

(12) Whatever place x I go, 2+2 = 4 in x;

The reason is that, if the hidden variable in (12) is assigned a value, different variable assignments do not produce a change in truth-conditions for the embedded phrase. Notably, the embedded sentence "2+2 = 4" has the same truth-conditions in "If I go to Paris, 2+2 = 4 in Paris" and in "If I go to London, 2+2 = 4 in London". These considerations tell us that there is no evidence that (11) should be interpreted as containing a binding relation at the semantic level. One should therefore refrain from drawing any conclusions as to the presence of a hidden variable at the level of syntax, which would allegedly result in a *reductio* of Stanley's binding argument.

To sum up, over-generation arguments are a threat for the Indexicalist only if they start from a genuine case of semantic binding. In these "good" cases, the proponent of the overgeneration challenge presses the Indexicalist to explain why his strategy demands that extra syntactic positions be countenanced, and hence that a new dimension of context-sensitivity should be admitted, even though no intuition of context-dependence strikes ordinary speakers[25].

---

[25]   Further references relevant to the discussion as to the legitimacy of alleged over-generation objections, are Stanley 2005, Cappelen and Hawthorne 2007, Zeman 2011.

6. Objection 3 – Indexicals, in what sense?

The Indexicalist position is based on the assumption that the elements that are articulated in the expressions' deep syntax function in a way relevantly similar to already recognised indexicals. With respect to some basic features of indexicals, this is very easy to see: on the assumption that a mark of indexicality is the fact that a certain expression varies its semantic value according to context, the indexical element x which is allegedly present in e.g. the logical form of "green" is such that it acquires a different semantic value in different contexts, thereby denoting a different property $green_1$, $green_2$, … $green_n$ in different contexts. This comparison is very straightforward and, at least in this respect, it equates genuine indexicals with the hidden indexicals postulated by Stanley. But what about what someone may indicate as *other essential features* of recognised indexicals? Do the hidden indexicals postulated by Stanley and others live up to the comparison with recognised indexicals even with respect to these other features?

Neale (2007) points out that another essential mark of indexicality is *perspective*. Words like "here"/"this" and "there"/"that" express a *proximal* and *distal spatial* perspective respectively; "now" and "then" express a proximal and distal *temporal* perspective respectively. Using an indexical like "I" signals what Neale calls the "perfectly proximal perspective" from which the speaker may be performing a certain predication—namely the *self* perspective; while using an indexical like "you" marks a "second-order perfectly proximal" perspective from which a certain predication could be effected—which we could cash out as "the most proximal perspective after the self perspective".

With respect to the aphonic, i.e. non-phonetically realised, indexicals posited by Stanley, Neale wonders if they can be described as having such a perspectival character, too. Take for example the aphonic indexical on location, call it *loc*, which Stanley postulates as figuring in the logical form of "to rain". Its function is that of referring to different objects in different contexts: but does it do so in a perspectival way? Neale observes that the perspective involved cannot be the same as that expressed by "here" or "there": for suppose I am in Paris, talking on the phone with Carlos, who is in the same room with Juan. Juan asks Carlos to ask me what the weather is in Paris. I utter "It's raining"; Carlos in turn utters "It's raining" to Juan. The hidden indexical *loc* cannot be equivalent to "here", otherwise Carlos could not have repeated my utterance the way he did; nor can it be equivalent to "there", for this would make my utterance absurd. So *loc* lacks an intrinsic perspectival character. Is then *loc* synonymous with some definite description or proper name, like "Paris" or "where Delia is"? Obviously not, because the sentence "It's raining" could have been used to say that it's raining in some place other than Paris, as well as in some place other than where I am at a certain moment.

Neale then summarises Stanley's theoretical claim as follows: Stanley wants *loc* to be an indexical, which is though *aphonic* and *aperspectival*. The only function performed by *loc* is that of referring to whatever location is salient at any stage of the conversation. But then, what is the communicative utility of such an item? Suppose *loc* had a phonic counterpart, call it *loke*. Anybody who wished to say that it's raining in some particular place—whether the place of the utterance or any other salient place—could utter "It's raining *loke*". But then, do we really need the contribution of *loke* in order to explain the fact that the proposition expressed by "It's raining" is about a certain particular place? Couldn't we just say, for example, that "to rain" *lexically* incorporates locations as playing a constant role in the interpretation of any occurrence of "rain"? This would be enough to explain why an occurrence of "It's raining" demands provision of a place, in order to express a proposition (see Neale 2007 and Taylor 2001 for this suggestion). Here is another reason to regard *loke* as syntactically superfluous: suppose that someone utters "It's raining here in Paris". If one is talking about Paris and one is indeed in Paris, then by uttering "here" and "in Paris" one is referring to the same place, even though with two different "modes of presentation": "here" presents Paris as the place of the utterance, while "in Paris" presents the place of the utterance as being Paris. Suppose

one now added to the previous sentence an occurrence of *loke*: "It's raining *loke* here in Paris". What is the contribution of *loke* to the information conveyed by the utterance? Neale's answer is: none. There would be no informational loss if one dropped the occurrence of *loke*, because *loke* has no communicative role: it just refers to whatever one wishes it to refer. From a methodological point of view, then, there is no good enough reason to recognise *loke* as expression with a *bona fide* role on one's syntax. As Neale points out

> *ceteris paribus*, we should posit such things, as far as is possible, only when there are strong empirical reasons for doing so, when doing so appears empirically unavoidable or narrows other options. If we are going to posit expressions with phonic properties but no communicative utility [like e.g. *loke*], this should be because we have been led to them by strong empirical considerations, not simply because the idea of such expressions is not itself incoherent. (Neale 2007: 67)

Now remember that, in Stanley's discussion, no such thing as a phonic location parameter is at stake, for all that Stanley discusses is the presence of an *aphonic* locational parameter, *loc*. So let's go back to *loc*: since *loc* has no phonic realisation and no communicative role, is there any reason for positing it besides that its presence is not incoherent with the rest of our syntactic posits? It seems that Stanley should provide such reasons, for otherwise one could observe, as Neale does:

> The idea of expressions with phonic properties but no communicative utility is not hard to get one's mind around[26]. Nor is the idea of expressions with communicative utility but no phonic properties[27]. But the idea of expressions with neither communicative utility nor phonic properties? The idea is strained. (Neale 2007: 67)

There are two other characteristic features of indexicals, highlighted by Cappelen and Lepore (2002), which Stanley's indexicals seem to lack. The first one is that *indexicals can enter in anaphoric relations*: that is, their reference can be picked up by a pronominal element that occurs after them in a sentence. For example, consider an utterance of "That is a duck and *it* is my favourite animal". Here the pronoun "it" is anaphoric for "that". Now Cappelen and Lepore surmise that, if hidden indexicals are genuine indexicals, they will be able to generate anaphoric relations, too. Consider a sentence like "It's raining". The real structure of the sentence is, by hypothesis, such that it contains a variable for locations. If this is the case, then the indexical element *l* should be able to generate anaphoric readings. But this is not the case. The sentence "It's raining and it is a big location" has no anaphoric reading, even though Stanley's theory predicts so. So the Indexicalist view fails to satisfy the prediction that hidden indexicals, just like overt indexicals, should be able to generate anaphoras.

The second characteristic feature of indexicals is that they give rise to *a priori* knowable truths. For example, the sentence "I am the person who is uttering this sentence" is going to be *a priori* true. The character of "I" is such that it imposes on any occurrence of "I" the requirement that its referent be the speaker of the utterance: Only knowledge of the word's character is required in order to derive that the sentence "I am the person who is speaking" is true (indeed, logically true). Now Cappelen and Lepore surmise that, if hidden indexicals are indexicals, then they will give rise to *a priori* truths, too. They therefore consider an occurrence of "Everyone": According to Stanley and Stanley and Szabó (2000), quantifier expressions contain a hidden indexical ranging over domains, so that the real logical form of the phrase is "Everyone *in domain D*", where D is to be saturated with a contextually salient domain. If things are like Stanley and Szabó claim, then the following sentence should be a priori true: "Everyone is in the contextually salient domain". This sentence

---

[26]   Think for example of the "it" in "It's raining". It has no communicative role, but it has phonic realisation.
[27]   Think for example of the aphonic PRO element in "Susan wanted PRO to sing".

should be known as true only in virtue of the fact that "everyone" is associated with an element ranging on contextually salient domains. However, this sentence is not an *a priori* truth, because it doesn't seem part of the meaning of "everyone" that it (even implicitly) points at any contextually salient domain. If this is the case, then Stanley and Szabó should at least "defend their departure" (Cappelen & Lepore 2002: 279) from the predictions released by the paradigm account on the semantics of genuine indexicals.

7. Objection 4 – Wrong Methodology

A very important issue is whether Stanley is complying with the correct methodology that one should adopt while investigating syntax. Stanley is pretty clear that the view he is advocating should apply to the *logical form* of sentences, where by logical form Stanley means the "real structure" that belongs to these sentences. His investigation is aimed not at a *revision*, effected in semi-formal English, of sentences of informal natural language, but rather at the *discovery* of what the true structure of certain linguistic constructions is. In other words, Stanley is not looking for a *regimentation* of certain linguistic constructions for the purpose of avoiding ambiguity and pursuing clarity and rigour; rather, he is looking for a *description* of the underlying structural features of sentences such that it can explain a certain range of phenomena.

In light of these methodological commitments, some theorists have wondered whether Stanley's conclusions are entirely motivated. The reason is that Stanley's postulation of hidden indexical elements in the logical form of e.g. "It's raining" stems from semantic considerations. However, is it methodologically correct to postulate syntactic elements just because this yields certain correct semantic results?

According to Stanley, the following is a methodologically sound principle: a syntactic item is correctly postulated whenever the presence of that item explains a behaviour which would remain unexplained in case that item was absent (see Stanley 2002: 152). In order for this methodological statement to be correct by the lights of the syntactician, it however has to be ascertained whether the behaviour to be explained is a semantic or syntactic behaviour. Explaining a syntactic behaviour (or misbehaviour) by postulating some hidden syntactic element seems an optimal way of applying Stanley's methodological statement. Let us review a virtuous way of applying this methodology. The example involves the motivation that has lead theorists to postulate the existence of what is known as a "PRO-element".

To briefly introduce the notion of PRO: a PRO element is an unpronounced item which is typically postulated in order to explain the grammaticality of certain linguistic forms. The clearest example of a well-motivated appeal to PRO is provided by those forms which feature an infinitive clause. Consider the sentence "John wanted to sing". The superficial syntactic structure of the sentence delivers a non-well-formed formula. That is because the infinitival clause "to sing", which is embedded under the main clause "John wants" doesn't exhibit a NP-VP structure, in that it only contains a verb-phrase node *<to sing>* which has no noun-phrase as its antecedent. This is problematic, because a failure to specify an antecedent for the verb-phrase *<to sing>* implies disregarding the thematic (or theta-) structure of the verb *to sing*, which necessarily requires an "agent" that does the singing. How is the theta-structure of the verb upheld, and hence the good formation of the sentence restored? Here syntacticians have resorted to the postulation of a hidden pronominal element PRO. The "real" structure of the sentence is therefore: *<John wanted PRO to sing>*, where the PRO element is a silent pronominal which inherits the reference of the proper name "John", which is in subject position in the upper-level structure of the sentence. Here the methodology is virtuously applied, given that both the *explanans* and the *explanandum* belong to the domain of syntax: since infinitival clauses show a deviant syntactic behaviour, a syntactic element is invoked to explain such a phenomenon.

After having made this detour into syntax, and namely into an example of well-grounded postulation of a syntactic element, let us now turn to the issue whether Stanley's hidden pronominal

elements are justified by the lights of syntactic methodology. On this issue, more than one author has advanced serious, if not fatal worries against Stanley's strategy.

Once again Neale (2007) argues that Stanley overstates his case. In an example like "Everywhere I go, it rains", there are *obviously* semantic or interpretive reasons to believe that the rain-location varies with the location in which the speaker goes. But this all, Neale emphasises, has to do with how we read or interpret the sentence, *not* with an account of the sentence's syntax.

> Stanley claims to have syntactic evidence for the claim that ["Everywhere I went, it rained]'s LF contains an aphonic variable *loc* bound by 'everywhere I went'. But in fact the evidence offered is *purely interpretive*: the mere existence of the undisputed, relativized interpretation of [the sentence]! Surely no one is going to argue with the well-known *interpretive* fact that [that sentence] has a relativized interpretation. The question is whether anything of *syntactic* significance follows immediately from this. Stanley assumes without argument that it does. (Neale 2007: 73)

Neale's point is that Stanley fails to distinguish between (i) relativisation, which is an interpretive aspect of sentences, such as that which obtains every time the values of a variable vary with those of a binding operator (e.g. a quantifier), and (ii) binding, which responds to syntactic motivations, e.g. well-formedness, coherence with other syntactic phenomena etc. The two aspects may come together, but they do not need to. On the one hand, variable-binding is not sufficient to get a relativised reading, because we can have cases of bound variables which do not determine relativised readings. For example, in "Ringo loves himself", we can say that "himself" is bound by the proper name "Ringo"; but the reading is not relativised, in the sense that there is no co-variation of the variable with respect to the values of any other element—in fact, we have a case of co-reference between the proper name "Ringo" and the pronoun "himself". On the other hand, variable-binding is not necessary for relativisation, in the sense that one can have relativised readings without having bound variables. An example of relativisation without syntactic binding is found in donkey anaphora, in which the relativised element (a pronoun, a variable) is in another sentence altogether with respect to the quantifier, and hence falls outside its scope, with the result that the quantifier cannot bind it. An example could be the sentence "John bought exactly one donkey and Paul fed it" (see Neale 2007: 73-8).

A criticism similar to the one moved by Neale is raised by Pupa and Troseth (2011). In their view, Stanley fails to comply with the syntactician's methodology, in that he offers no reasons, beyond the interpretive ones, to think that sentences contain aphonic elements. As opposed to this, orthodox syntactic analysis proceeds mainly from (i) well-formedness considerations; (ii) distributional considerations, i.e. two elements are said to be syntactically incompatible if they compete for the same structural position; (iii) cross-linguistic considerations: for example, if a certain element is phonetically empty in a certain language, but it is phonetically realised in another language, that is a reason for postulating it. Stanley's arguments seem to fly in the face of this methodology. Here is how the authors express their concerns:

> From a syntactic perspective, we find the binding argument peculiar. It never invokes any well-formedness considerations at all. And while the syntactic analysis of a sentence may ultimately have interpretational effects, one shouldn't resolve interpretational puzzles by positing syntactic constituents without developing a well-formedness argument in their favour. The binding argument, however, makes no such gestures. As such, it's difficult to imagine the binding argument grounding any syntactic thesis. This we note as a theoretical challenge for the binding argument. (Pupa and Troseth 2011: 194).

Pupa and Troseth also echo Neale's objection from the conflation of relativisation and binding when they point out that semantic binding and syntactic binding need not coincide: that is, the fact that the semantic value of an element like a pronoun or variable is controlled by an operator like a quantifier may not imply any form of syntactic connection between the elements. Take for example the sentence "In most of Jennifer's classes, she fails exactly three Frenchmen". Stanley and Szabó (2000) argue that its logical form should be such that the noun "Frenchmen" contains an aphonic domain restrictor $<g>$, which is controlled by the upper quantified phrase "In most of Jennifer's classes". The result is

(13) In most of Jennifer's$_1$ classes x, she$_1$ fails exactly three Frenchmen$_{<in\ x>}$

The problem with this way of cashing out the sentence's logical form is that the quantifier phrase "most of Jennifer's classes" cannot bind the restrictor $<g>$ in "Frenchmen" because, as the authors observe, "g is a sub-syntactic element. Accordingly, it is not the type of element that can be c-commanded. If 'most' does not c-command g, it cannot bind g." (Pupa and Troseth 2011: 194). There are other cases in which we have a semantic binding without a syntactic binding. Take for example the sentence

(14) Some manager didn't manage to deceive every shareholder in the company

In (14), a bound reading is easy to obtain to the extent that it is easy to read "Some manager" as being some manager *in the company*. Does this mean that the quantified phrase "every shareholder in the company" binds an aphonic domain restrictor $<g>$ associated with "some manager"? According to syntax, this is impossible because "restrictions on syntactic movement prohibit the lower DP 'every shareholder in the company' from taking a scope position higher than the negation 'didn't'. Since the lower DP cannot, in virtue of the interaction between syntactic movement and negation placement, out-scope 'didn't', and thereby out-scope 'some manager', semantic and syntactic binding cannot coincide on threat of ungrammaticality." (Pupa and Troseth 2011: 194-5). In other words, since the the determiner-phrase (DP) "every shareholder in the company" is located in a *lower* hierarchical position than the DP "some manager", and since it cannot be raised so as to become hierarchically higher than it, it cannot bind it.

Further considerations, which yield to a demotion of Stanley's strategy in general, are forcefully put forward by Collins (2007). Collins focusses on Stanley's key cases of quantifier domain restriction (see Stanley 2000). In short, Stanley takes it that a quantified phrase such as "every question" in the sentence "Every student answered every question" must contain a hidden variable in its logical form, thus having the form $<every\ question_{(x)}>$; if this were not so, the fact that the domain of the second quantified phrase ("every question") is bound to the domain of the upper-level quantified phrase "Every student" cannot be explained.

Collins wonders what kind of syntactic relation could realise the semantic binding between the first and the second quantifier. He offers some answers, which he shows to all be inadequate. A first potential syntactic relation could be *c-command*, (where a node α c-commands another node β in a parse-tree if α and β are not hierarchically related, i.e. they do not "dominate" each other, and the first branching that dominates α also dominates β). However, c-command already obtains anyway between "every student" and "every question", and it in no way reveals whether any bound variable is there to be c-commanded as well. A second potential syntactic relation could be *adjunction* (where adjunction is a relation that obtains between a node α and a node β when node β is combined with node α without this causing any change in the category of α, for instance the phrase category). This however, is not plausible either, because adjuncts are generally phrases, and it is difficult to see covert variables as equivalent to phrases. Moreover, adjuncts are essentially optional—their specification is not demanded at the level of theta-structure—so their nature is deeply different from

the nature of the alleged hidden variables, which occupy positions that are to be filled in a mandatory way.

Drawing from these considerations, Collins concludes that the bound readings which should provide evidence of the articulation of hidden variables in logical form are not supported by any serious syntactic relation. Syntax offers too little support to Stanley's binding argument, for it's not clear which syntactic relation should realise the semantic relation of binding which Stanley notes. As Collins has it, the implicit premise that syntax can be derived from semantic considerations is not adequately supported:

> Overall, Stanley's argument here, as far as I can see, is premised upon the silent assumption that the nature of syntax should be read off our stable semantic intuitions. The specific issue of binding is a red herring in so far as the examples offered do not in fact exhibit syntactic binding […]; rather, it is particular readings that purport to exhibit binding; and even if the readings are correct, nothing would ipso facto follow as regards syntax save for the c-command relation, which is already in place.

To sum up, the concern raised by Neale, Pupa & Troseth and Collins is about the *methodological orthodoxy* of Stanley's arguments. These authors do not question either the logic or the dialectic of Stanley's defence; they simply wonder whether the whole project is enough supported by syntactic arguments, and after a thorough analysis, they conclude it isn't.

8. Zero-Readings

In this section, I would like to call attention on a phenomenon that I shall call "Zero Reading". I shall describe the phenomenon, make some examples and finally argue that a Stanley-style solution cannot capture it. This will count as a fault for all those accounts inspired by Indexical Contextualism, and shall prompt us to search for an alternative solution.

First of all, some simple facts about the way we talk about rain, being ready, being green. Not every time we talk about rain do we talk about a place where it rains. Sometimes we are just interested in rain as *a phenomenon of rain-falling*. Analogously, not every time we talk about someone or something being ready do we talk about a purpose for which someone or something is ready: sometimes we may just be interested in *the state of being ready*. Similarly, not every talk about something being green is talk about a way, or part in which something is green. It seems that speakers may also be interested in the simple *state of being green*. Note that this *doesn't imply* that, when we are engaged in this kind of talk, we somehow believe or presuppose that rain-falling events do not occur at a place and a time; or that states of being ready are not states of being ready for some activity; or that states of being green do not hold in certain respects, or part. I think we can safely say that our *concept* of rain-falling events continues to be one in which these episodes happen at a location and a time, and analogous considerations apply to states of being ready and being green. Hence, the structure of our concepts of rain events, readiness and greenness does not change, even though it seems that in talk we can "bracket", or "leave in the background" aspects of these events and states such as locations, purposes, parts; we can temporarily treat expressions like "to rain", "being ready", "being green" as if, at the syntactic-linguistic level, they were monadic predicates, even though at the conceptual/semantic level they are clearly more complex.

I shall call a *zero-reading* of "It's raining", "X is ready", "X is green" a reading of an utterance of these sentences in which the predicates "to rain", "being ready" and "being green" are felicitously used with a monadic syntactic structure. Here are some examples:

(15) Whenever the temperature and humidity in the air reach a certain level, it rains[28];

---

[28]   See also Cappelen and Hawthorne 2007 for a similar example.

(16) When it rains, caution in driving is advised.
(17) I hate it when it rains!

All these sentences speak about episodes of rain-falling as if they were self-contained, self-standing events. In order to see this, one could try to substitute every occurrence of "it rains" with a construction involving the phrase "an event of rain takes/took place". Thus, sentence (15) states that whenever the temperature and humidity in the air reach a certain level, *an event of rain takes place*. Here locational aspects play a role at the truth-conditional level—after all, that an event of raining takes place necessarily presupposes that there is a location—but this role need not correspond to anything at the syntactic-linguistic level. Sentence (16) is also paraphrasable in the following way: whenever an event of rain-falling takes place, caution in driving is advised. Here again locational aspects play a role at the truth-conditional level, even though they need not be articulated at the level of syntax. Finally, imagine a child who exclaims (17), after being told that his favourite fun park, where the whole family was planning to go, is currently under a horrible storm. The child's utterance could be read as a disapproval of events of rain in general, to the extent that they prevent him from engaging in entertaining activities. Again, the fact that every event of raining has a locational aspect is part of the truth-conditions, but it needs play no role at the linguistic-syntactic level.
Zero-readings can be extended also to other predicates, such as "ready":

(18) Whenever she is prepared for an activity, Jill is ready;
(19) John being on time usually depends on Jill's being ready;

I take sentences (18)-(19) to all have occurrences of the predicate "being ready" which are to be taken as occurrences of a *monadic* predicate. Like in the previous cases, we could easily see this by substituting every occurrence of "is ready" or "being ready" with the phrase "being in a state of readiness". Thus, sentence (18) states that if Jill is prepared for any activity, then she is in a *state of readiness*. That readiness is readiness with respect to an activity is certainly part of the truth-conditions, however, the predicate needs not be dyadic in its structure in order for such an aspect to be part of what makes an utterance of this sentence true or false. Consider now an utterance of (19): it can be viewed as stating that John's being on time usually depends on Jill's being in a state of readiness. The relation to an activity certainly plays a truth-conditional role, though nothing follows from this concerning the syntactic structure of the predicate.
Finally, consider the following group of cases involving the predicate "being green":

(20) Whenever something or some salient part of it looks like this, it is green;
(21) When number X appears on the display, then you are looking at something green;

Both uses of the expression "green" should be such that they speak about the state of being green, or the possession of the property of being green. Again both sentences could be paraphrased by substituting occurrences of "is green" with phrases like "being in a state of greenness" or "possessing the property of being green". Imagine that the speaker is teaching to a child the meaning of "green" by ostension. By pointing at a green patch, the speaker utters (20): the content of his utterance could be to the effect that whenever something or some salient part of it looks like the patch demonstrated, it has the property of being green. Similarly, suppose that a speaker utters (21) while illustrating the functioning of a wave-length detecting device to a colour-blind person. The content of the utterance could be rephrased as being to the effect that, whenever the device gives wave-length X as a response, the person is looking at something which is in a state of greenness. Again the metaphysical fact that greenness is greenness in some respect, or part, is part of the truth-conditions but it needs not be articulated at the syntactic level.

In this section, I have introduced the phenomenon of Zero Readings, i.e. readings in which we speak about raining, being ready, being green etc. while using the relevant predicates as monadic and not polyadic. In the next section I shall be concerned with the issue of what treatment an Indexical Contextualist should reserve for these readings.

9. Indexicalism on Zero Readings

How should the Indexicalist deal with Zero Readings? Remember that the Indexicalist claims that "to rain", "being ready", "being green" contain an argument-place in logical form for locations, activities and parts respectively. In none of the (15)-(21) examples there is a way of filling in that argument-place. This implies that the argument position ultimately remains vacant. This entails that, for example, once we assign the correct logical form to sentence (15), it should turn out that it is really an *open sentence*, as in (15'). And the same holds when we write down properly the syntactic structure of (18) and (20): what we discover is that, since these sentences contain an extra argument-place in logical form, their logical form is really that of *open* sentences, as one can see in (18') and (20').

(15') Whenever the temperature and humidity in the air reach a certain level, it rains*(x)*;
(18') Whenever she is prepared for an activity, Jill is ready*(x)*;
(20') Whenever something or some salient part of it looks like this, it is green*(x)*;

Typically, open sentences are true or false only with respect to an *assignment* of values to the free variables that occur in them. The variables that occupy the relevant argument-places therefore need to be assigned a value; the argument positions cannot stay vacant or unsaturated, on pains of the sentences not expressing any evaluable content. The Indexicalist is therefore bound to the claim that (15), (18) and (20), as they stand, are unevaluable until a value is assigned to the free (covert) variables that occupy argument positions in their logical form. This strikes me as simply not true: each of these sentences expresses an intelligible, evaluable content.

The Indexicalist could offer the following explanation of why (15), (18) and (20) express intelligible and evaluable contents. There are two possibilities: Either the variables are "freely" assigned a specific value contextually, or they get bound by a quantifier. Take for example sentence (15). The Indexicalist could say that the sentence either receives a "free" reading like that in (15a), where the unconstrained location variable gets assigned a value in context (e.g. Paris), or a "bound" reading like the one in (15b), in which the quantifier "whenever" binds a locational variable in the embedded clause "...it rains":

(15a) Whenever the temperature and humidity in the air reach a certain level, it rains *in Paris*;
(15b) For every time *t* and every place *l* in which the temperature and humidity in the air reach a certain level, it rains at *t* in *l*.

(15a) strikes me as utterly mistaken as a way of spelling out a supposed zero-reading. As far as (15b) is concerned, my inclination is to say that it is fine, but this doesn't nonetheless put the Indexicalist in a better position. This is because bound readings like (15b) are available even if one doesn't believe that "to rain" has a complex syntactic structure *rain(x, y)*. For example, one may believe that a variadic function expands the adicity of "to rain" in appropriate contexts, such as quantified ones; even a proponent of the event-analysis could devise such a bound reading, as Cappelen and Hawthorne (2007) show (see the following section for more details). What is central to the objection I am currently making is that Indexicalism cannot rule out a reading like (15a), on the account that the he believes that "to rain" has a dyadic syntactic structure *rain(x, y)*. This implies that the option is available to read (15) as an open sentence, which in turn entails that (15a) remains at the semanticist's disposal. Moreover, even if the Indexicalist can avoid (15a) by offering a bound

reading like (15b), still he cannot explain the fact that zero-readings seem to involve monadic predicates by saying that the predicates at issue are *indeed* monadic. As opposed to this, the non-Indexicalist can avoid regarding (15) as an open sentence, thereby avoiding also reading (15a), in virtue of the fact that, in her view, "to rain" has a monadic structure *rain(x).* Envisaging "to rain" as having a monadic structure blocks the reading of (15) as an open sentence, and hence blocks (15a). Moreover, it allows the theorist to explain the fact that zero-readings seem to involve monadic predicates with the claim that they *do*, after all, involve monadic predicates.

The same moves would apply to examples (18) and (20). Let us examine (18) first. Here again the Indexicalist has to explain that, despite the fact that the sentence is not strictly speaking evaluable—since it has the logical form of an open sentence—utterances of it express intelligible and truth-evaluable contents. There would be two ways of explaining this fact: either the variable *x* that occupies the argument-place for activities gets a "free" assignment, according to which activity is salient in context (for instance, the activity of *skiing*); or the variable *x* gets bound by the quantifier "whenever". The resulting readings are displayed in (18a) and (18b):

(18a) Whenever she is prepared for an activity, Jill is ready *for skiing*;
(18b) For every activity *x* she is ready for, Jill is ready for *x;*

Once again, (18a) looks just wrong as a way of making sense of a zero-reading, while (18b) looks fine. However, note that one can interpret (18b) as displaying a bound reading even if one doesn't believe that "being ready" has a complex syntactic structure like *ready(x, y)*. One could, for example, explain it by saying that the presence of the quantifier increases the predicate's adicity; even an event-semanticist could obtain the bound reading (see the following section for more details). So reading (18b) may be correct as a way of making sense of the fact that (18) is perfectly evaluable, but it is not exclusively available to the Indexicalist. What is important for the objection is that a reading like (18a) is in principle available for the proponent of Indexicalism: this seems to be the case, for Indexicalism has it that "being ready" has a dyadic syntactic structure *ready(x, y)*, an assumption which cannot in principle block a reading of (18) as an open sentence, and therefore cannot block reading (18a). By contrast, a non-Indexicalist view, which has it that "being ready" has the monadic structure *ready(x)*, can avoid reading (18) as an open sentence, and therefore can avoid reading (18a). A further advantage is that the Indexicalist, who maintains that "to rain" has a dyadic structure, can't explain the fact that zero-readings seem to involve monadic predicates with the idea that the predicates are *indeed* monadic, while such a straightforward explanation is available to the non-Indexicalist.

The same remarks hold for (20)-(21). The Indexicalist is forced to say that (20) has the logical form displayed in (20'), and therefore that it amounts to an open sentence whose evaluation is impossible until the free variable associated with "green" is assigned a value. This clashes with the fact that the sentence expresses an entirely evaluable content. The explanation that the Indexicalist could invoke would be analogous to the one put forth with relation to (15) and (18). Either the variable gets a "free" assignment, in accordance with the particular way of being green that is salient in context, the result being something like (20a); or the variable gets bound by a quantifier, the result being (20b):

(20a) Whenever something or some salient part of it looks like this, it is green *on the outer surface;*
(20b) For every object *x*, for every part *y* which looks like this, *x* is green with respect to *y.*

Once again, it seems that (20a) doesn't capture the zero-reading of (20). (20b) seems fine, but again note that the bound reading exhibited by (20b) is available at the interpretive level even if one doesn't assign the predicate "being green" a complex structure like *green(x, y).* For instance, one could explain the bound reading via the effect of a variadic function which increases the adicity of

the predicate; one could even adopt an event analysis in which the respect variable y gets bound, yet it doesn't occupy an argument-place position, but rather an adjunct-position (more on this in the following section). What is truly problematic for the Indexicalist is the fact that (20a) is available, which seems to be the case given that Indexicalism regards the predicate "being green" as having a dyadic syntactic structure *green(x, y)*. This causes the Indexicaist to interpret (20) as an open sentence, thus allowing for a reading like (20a). By contrast, the non-Indexicalist view claims that "being green" has a simple monadic structure *green(x)*. This blocks the reading of (20) as an open sentence and, ultimately, blocks also reading (20a). Furthermore, the non-Indexicalist has a direct explanation for zero readings—namely that they seem to involve monadic predicates because they, as a matter of fact, *do* involve monadic predicates; the Indexicalist, by contrast doesn't have access to this explanation.

A qualification is called for: of course I don't want to deny that "free assignment" readings (15a), (18a) and (20a) are *possible* and indeed *perfectly admissible* readings of (15)-(18)-(20), in appropriate occasions. One can imagine a scenario in which a speaker utters any of these sentences in order to communicate exactly what (15a), (18a), (20a) express. Note, however, that (15a), (18a), (20a) too are available also to the non-Indexicalist. For example, one could account for them by invoking the effect of variadic functions, which increase the predicates' adicity whenever a potential argument-filler (Paris, skiing, etc.) is salient. Be that as it may, the problem is that (15a), (18a), (20a) are inadequate as far as spelling out zero-readings is concerned. My objection that they are troubling for the proponent of Indexicalism is therefore relative to the adequacy of (15a), (18a), (20a) *as ways of rendering zero-readings*, not as ways of rendering *merely possible readings* for (15), (18), (20) in contexts where this is appropriate.

The final result is that the Indexicalist account cannot capture zero readings, even if these readings appear as available in natural language. The reason is that, in the Indexicalist's view, predicates such as "to rain", "being ready", "being green" are all associated with extra argument-places in logical form. This makes it that sentences like (15)-(21) turn out as *open sentences*, something which goes counter to our intuitions. The Indexicalist could try to explain zero readings by resorting to *bound readings* like the ones in (15b), (18b) and (20b). These readings, however, are available also to the non-Indexicalist. Ultimately what should trouble the Indexicalist is the fact that readings in which the alleged free variable is assigned a value contextually are always available to his view—like (15a), (18a) and (20a). These readings cannot be ruled out by the Indexicalist, while they can be ruled out by the theorist who believes that "to rain", "being ready", "being green" exhibit a monadic structure. Ultimately then, if zero-readings stem from our seeming use of "to rain", "being ready", etc. as monadic predicates, a non-Indexicalist can explain these readings with the idea that we are *indeed* using monadic predicates *rain(x), ready(x)*, etc. The Indexicalist, on the other hand, has to resort to bound readings in order to explain the fact that we seem to use monadic predicates—when in fact, in his view, we are using dyadic predicates *rain(x, y), ready(x, y)*, etc.

10. Davidson-Inspired Accounts

In this section, I'll sketch how a Davidson-style analysis in terms of events allows to capture zero-readings. I'll also argue that, if such an analysis manages to accommodate *also* bound readings, then the overall account is to be preferred to a Stanley-style approach. First, though, I shall spend some words of introduction to the fundamentals of a Davidsonian event-semantics.

In his famous paper "The Logical Form of Action Sentences", Donald Davidson puts forward an analysis of action sentences in terms of events. His starting point is that of accounting for two problems associated with this kind of sentences. The first problem may be called the *polyadicity problem*, and it concerns the fact that, in describing an action, an indefinite number of details could be added by making either adverbs or prepositional phrases follow an action verb. For example, one can go from sentence (22) to sentences (23) and (24) by simply attaching the words "slowly" and "at midnight" to the sentence:

(22) John buttered the toast;
(23) John buttered the toast slowly;
(24) John buttered the toast slowly at midnight.

The fact that one can get sentences (23) and (24) from (22) may lead one to think that one is introducing a new action verb with an *increased adicity* every time one specifies a detail of the action by means of either an adverb ("slowly") or a prepositional phrase ("at midnight"). This, however, can't be right for many reasons (see Davidson 1967: 83-4). The second problem, which we may call the *equivalence problem,* corresponds to the need to license the derivation of different but logically equivalent descriptions of the same action. It seems desirable to license inferences of the following form: if it is the case that I drove my spaceship to the Morning Star, and the Morning Star is the Evening Star, then it is the case that I drove my spaceship to the Evening Star.

Davidson's basic idea is that verbs of action are predicates of events. So for example, the verb "to kick" is thought of as having logical form *kick(e)*, and its denotation is the set of events that are kicking-events. This conception imposes a change in the way the logical form of action sentences is represented. The change consists in conceiving action sentences as taking the form of *existential quantifications over events*. Thus, in Davidson's analysis, the logical form of "Shem kicked Shaun" is the following:

(25) (Ex)(kicked (Shem, Shaun, x))

To be read as "There is an event x such that x is a kicking of Shaun by Shem". This solution allows the theorist to solve both the polyadicity and the equivalence problem. With relation to the polyadicity problem, it allows to say that the expanded sentences obtained from the addition of adverbial expressions, like (23) and (24), do not involve the introduction of a new verb with increased adicity. This follows from the fact that, in Davidson's framework, the addition of an adverbial expression is tantamount to the *introduction of a new conjunct* and not to the creation of a new slot in argument structure. This can be seen once one re-analyses a sentence like (24), whose logical form now becomes

(24') (Ex) buttered(John, the toast, x) & slow(x) & at midnight(x).

To be read as: "There is an event x such that x is a buttering of the toast by John *and* x is slow *and* x is at midnight", where adverbial modifiers such as "slowly" and "at midnight" are treated as conjuncts. This analysis proves definitely more perspicuous; one remarkable feature is that it explains the fact that (24) entails (23) and (22) simply by appealing to conjunction elimination. This explanation connects to the solution to the polyadicity problem, in that it allows one to regard complex action sentences such as (22)-(24) not as introducing new verbs with higher adicity, but as containing the same verb with the same adicity, only followed by an extra conjunct. With relation to the equivalence problem, the account licences inferences like the Morning Star/Evening Star inference simply in virtue of the substitutivity of coreferential terms.

As one may expect, the event analysis can be applied also to sentences that do not contain action predicates. As Davidson himself writes:

> In general, what kinds of predicates do have event-places? Without pursuing this question very far, I think it is evident that if action predicates do, many other predicates that have little relation to action do. (Davidson 1967: 93)

Action verbs generally suggest both agency and intentionality, i.e. they suggest that the event

corresponding to the action was brought about by an "author" and that the author had the intention to perform the action. Verbs that have nothing or little to do with action are verbs that describe exactly "events" in the sense of "happenings" ("The bottle fell from the shelf") but also states ("The room is dark") and processes ("The pot is heating")[29]. An event analysis can be applied to all of these groups of verbs, as in Parsons's 1990 *Events in the Semantics of English*, which provides a development and refinement of Davidson's semantics.

We spoke of zero-readings of "It's raining", "Jill is ready" and "The leaves are green" as involving the rather informal notions of *events* of rain-falling, *states* of readiness, *states* of greenness. This terminology was helpful because it allowed one to see how language users speak about rainings, readiness and greenness as if they involved monadic predicates denoting monadic properties. This informal terminology can now be made to fit into the formal framework of event-analysis. I therefore propose to adopt an event-analysis for those verbs which can generate zero-readings; these readings shall be rendered as existential quantifications over an event-variable. So for example, zero-readings of "It rains" shall receive the logical form in (26), to be read as "There is an event x such that x is a rain-falling"; zero-readings of "Jill is ready" shall receive the logical form in (27), to be read as "There is an event x such that x is a being ready of Jill"; and zero-readings of "The leaves are green" shall receive the logical form in (28), to be read as "There is an event x such that x is a being green of the leaves":

(26) Ex (rain-falling, x);
(27) Ex (Jill, ready, x);
(28) Ex (green, the leaves, x);

(The representations in (26)-(28) follow the Davidsonian proposal for the sake of simplicity, but of course one could choose a more fine-grained representation such as Parson's, in which thematic roles (such as "AGENT", "THEME", "EXPERIENCER" etc.) are written as separate conjuncts). What is of interest for present purposes is that the essential feature of the Davidsonian account, i.e. the idea of *quantification over an event-variable*, is exactly what is needed in order to account for zero-readings. Other authors such as Cappelen and Lepore (2002), Cappelen and Hawthorne (2007) and Recanati (2010) already adopted an event analysis for sentences such as "It's raining", in order to account for a range of interpretive phenomena that Indexicalism has trouble in capturing. I am very much sympathetic to the strategy adopted by these authors, and I am therefore happy to embrace the same analysis; the significance of this move for the purposes of this investigation will appear more clearly in **section 11**, where I shall draw a methodological moral which will hopefully cast new light on the event-semantics framework and will strengthen the motivation for adopting it.

At this point, the challenge for anyone who opts for an event-analysis is to guarantee that it not only explains those phenomena, such as zero-readings, that Indexicalism cannot capture, but also that it offers the same benefits originally advertised by the proponents of Indexicalism. Since an account for the phenomenon of binding was among the advantages originally afforded by this theory, the question becomes: can an event-analysis explain binding? Authors such as Cappelen and Hawthorne (2007) have accepted this challenge. Very roughly, Cappelen and Hawthorne's solution consists in analysing sentences like (29) as having the logical form in (30):

(29) Every time John lights a cigarette, it rains;
(30) For every time t of John's lighting a cigarette, there is an event e of rain-falling$_{f(t)}$ occurring at t' and $t' = t$.

---

[29] There are at least two categories that one may place under the general heading of "mere happenings". On the basis of Vendler (1957), one could distinguish between *accomplishments*, i.e. events that require a previous stage of elaboration and have a culmination (making a cake, washing the car) and achievements, i.e. events that are instantaneous (reaching the summit, arriving, delivering).

The bound reading is explained in the following way: the temporal quantification expressed by the phrase "Every time John lights a cigarette" acts as a *domain-restrictor* for the existential quantifier "there is an event of rain-falling". As one can see, the existential quantification in (29) is restricted through a function $f(t)$ from the times t of cigarette-lighting to the events e of rain-falling. Binding as a semantic phenomenon is therefore turned, at the level of semantic analysis, into a phenomenon of restriction of domain of quantification. Cappelen and Hawthorne advocate more than one reason for recommending this analysis of binding, as contrasted with Stanley's Indexicalist approach.

An aspect in which Cappelen and Hawthorne's account does better than Stanley's is the existentially quantified versions of (29), like for example "Whenever someone lights a cigarette, it rains". In order for the bound reading to obtain here, a variable should be bound not only by the temporal quantifier "whenever", but also by the quantifier over individuals "someone". This cannot be done in Stanley's account, for the verb "to rain" is only associated with a temporal variable t. On Cappelen and Hawthorne's account, however, the problem is easily solved, because there is a way of representing individuals *via* variables in the sentence's semantic analysis. The result is (31), where use of variables for individuals is in boldface:

(31) For all times t, for all events e, **for all people p**, if e is a smoking by **p** at t, there is an event at t that is a raining $_{f(e)}$.

Consider now the following example:

(32) Whenever the river overflows, there is panic.

Here we want a bound reading to tell us that for all the times t the river overflows, there is panic at t in the place where the river overflows. We also want the binding to track the fact that the river can overflow in more than one place at the same time, and so that there can be panic in more than one place. This cannot be done if, following Stanley, we postulate a function from times t when the river overflows to the place where the over-flowing happens at t, because there would be no way for the function to "bring us" to different places: what we need is a way of quantifying over episodes of over-flowing, such that (32) can be made true by more than one overflowing at the same time t (at different places). Cappelen and Hawthorne's account delivers (33):

(33) For all times t and all events e, if e is an overflowing of the river at t, there is an event that is a panicking $_{f(e)}$ at t.

Where $f(e)$ is precisely the function we wanted from events of overflowing (and the related places) to events of panicking.

In this section we have seen the potential advantages of an event-analysis of predicates such as "to rain"; this kind of analysis (a) captures zero-readings, i.e. readings in which predicates such as "to rain", "being ready", "being green" seem to be treated as monadic; and (b) explains the phenomenon of binding at least as well as (if not better than) Indexicalism. What remains to be ascertained, in light of the methodological concerns raised in **section 7** of this chapter, is whether an event-analysis can be considered a *bona fide* syntactic theory, or if it is to play any different role.

11. Philosophical Analysis or Empirical Hypothesis?

Although I do agree with its proponents that an event-analysis may be superior to an Indexicalist hypothesis, I do have some remarks as to what these authors take themselves to do and as to what they are ultimately "allowed" to do, philosophically speaking. It could be argued that, in light of the objections raised by Neale, Collins and Pupa and Troseth, if it is true that the Indexicalist's proposal

cannot be taken seriously as a syntactic option, and if a Davidson-style account is put forth as an alternative to Indexicalism, then it cannot be taken seriously as a syntactic account *either*. The best it can achieve is offering a *logico-philosophical analysis* of a range of phenomena that have to do with how language users interpret certain sentences. My aim in this section is to scrutinise the consequences of this line of argument.

To start with, we need to distinguish between two different enterprises one can engage in, while trying to explain, capture, clarify or simply make sense of a semantic phenomenon: on the one hand, one can choose to engage in the task of providing a logico-philosophical analysis of the phenomenon. This may imply using the resources of a formal language in order to isolate the structure of a certain range of linguistic constructions and eventually clear up the confusion and ambiguity generated by natural language. If the problem is, for example, a confusion concerning the truth-value of sentences like "The king of France is bald", then the Russellian theory of descriptions is an example of how one can employ the resources of a formal language like that of first-order logic with predicate calculus in order to give a determinate answer to the question.

Another road one could take is that of working out a theory of syntax for a certain portion of natural language. The enterprise here is radically different in that it implies delivering an *empirical hypothesis* as to what the real structure is of a certain kind of linguistic construction; any claim in this field needs therefore further justification, based on observable, related syntactic phenomena that can act as evidence. Suppose, for example, we need to explain the ambiguity of the following sentence:

(34) Visiting relatives can be boring.

Here we know that there is a lexical ambiguity, in that the word "visiting" can be interpreted both as the continuous form of the verb "to visit" and as an adjectival modifier. This, however, cannot be all there is to the ambiguity of (34); (34) is not entirely analogous to a typical case of lexical ambiguity like "John went to the bank". The reason is that, depending on the interpretation of the word "visiting", the role of the word "relatives" in the sentence changes. If "visiting" is to be read as the present continuous of "visit", then the word "relatives" results the object of the predicate *to visit*; if "visiting" is to be read as an adjective, the word "relatives" no longer satisfies this role, but simply acts as the modified subject of the whole sentence. This connection between the interpretation of the lexical items and the structure of the sentence is not what makes the word "bank" ambiguous, for interpreting the word "bank" as "river side" or "financial institution" leaves the structure of the sentence altogether unmodified. We therefore need to account for the ambiguity of the sentence in terms different from the lexical ones. Here a genuine syntactic assumption is called for, in the sense that a new syntactic category besides the lexical one has to be countenanced in order to correctly describe the ambiguity at interest. The category of "phrase" becomes useful here: Using the "phrase" category allows to correctly locate the ambiguity at the level of the complex "visiting relatives", rather than at the level of the simple word "visiting", and therefore allows to accurately set apart the structures of the two distinct sentences that could be expressed by an occurrence of (34). We can spell out the structure of the sentence under the first reading as in (34a), while in the case of the second reading, the structure can be represented as in (34b):

(34a) [[[Visiting]$_{VP}$ [relatives]$_{NP}$]$_{NP}$ [can be boring]$_{VP}$]$_S$
(34b) [[[Visiting]$_{AP}$ [relatives]$_{NP}$]$_{NP}$ [can be boring]$_{VP}$]$_S$

The differences between the first and the second reading concern the structure of the phrase "Visiting relatives", and not just the lexical category of the word "visiting". In (34a), the occurrence of "visiting" constitutes a verb phrase and since the verb "to visit" is transitive, the occurrence of "relatives" fulfils the object role; in (34b), the occurrence of "visiting" constitutes an adjectival

phrase and, an adjective being a modifier, it calls for a modifiee—in this case, the occurrence of "relatives". Postulating phrases is a substantial syntactic move, which needs further justification on purely syntactic grounds (see Chomsky 1957, Radford 1981, 1988). For the present purposes, the example is just aimed at giving an illustration of what it means to take a syntactic route in order to explain a semantic phenomenon.

Stanley himself in his "Context and Logical Form" somewhat indirectly mentions the difference between a purely logico-philosophical and syntactic task, by distinguishing two different senses of the notion of "logical form":

> [T]here is certainly no one uniform use of the expression "logical form". But there are two distinguishable senses underlying its many differing usages. [...] Perhaps the most prevalent tradition of usage of the expression "logical form" in philosophy is to express what one might call the *revisionary* conception of logical form. According to the revisionary conception, natural language is defective in some fundamental way. Appeals to logical form are appeals to a kind of linguistic representation which is intended to replace natural language for the purposes of scientific or mathematical investigation. Different purposes may then give rise to different regimentations of natural language. [...] According to the second tradition of usage, which one might call the *descriptive* conception of logical form, the logical form of a sentence is something like the 'real structure' of that sentence (e.g., Harman (1972)). On this approach, we may discover that the 'real' structure of a natural language sentence is in fact quite distinct from its surface grammatical form. Talk of logical form in this sense involves attributing hidden complexity to sentences of natural language, complexity which is ultimately revealed by empirical inquiry. (Stanley 2000: 391-2)

As I understand Stanley, the notion of revisionary conception of logical form can be related to what I have called a logico-philosophical analysis, while the notion of descriptive conception of logical form can be related to the task of working out an empirical hypothesis as to the real syntax that underlies certain linguistic expressions. It is important to emphasise that, in his article, Stanley pursues exactly the second, descriptive sense of "logical form" in trying to give an account of the way pragmatic effects are traceable to logical form. As he himself writes:

> To say that all context-dependence is traceable to logical form in a revisionary sense of "logical form" might be taken to be the trivial claim that, for purposes of interpretation, one should replace natural language by a notation in which all context-dependence is made explicit in the favored notation. It is not in this sense that I intend the thesis. [...] It is in [the descriptive] sense that I intend the thesis that all context-dependence is traceable to logical form. What I shall defend is the claim that all truth-conditional context-dependence results from fixing the values of contextually sensitive elements *in the real structure* of natural language sentences. (Stanley 2000: 391-2, my italics).

Stanley's arguments and, in general, the arguments offered by proponents of Indexicalism are therefore to be understood as supporting claims at the level of *syntactic theory*. However, a certain number of theorists (Neale, Collins and Pupa and Troseth) objected that these arguments fail to give the appropriate support that any thesis in the theory of syntax calls for: the arguments set forth by these authors were surveyed in **section 7** of the present chapter. Now, one may wonder whether the same objections could be raised against the event-based accounts advocated by Cappelen & Lepore, Cappelen & Hawthorne and Recanati. I suspect that these approaches would fare no better than Indexicalism, if their methodological pedigree was tested by a syntactician: just like Indexicalism,

Davidsonian approaches posit entities such as hidden event-variables in order to explain a number of *logical and semantic* relations. This, as we have seen, is methodologically insufficient in order to license any claim as to the real syntactic structure of sentences. The upshot is that opponents of Indexicalism, just like Indexicalism itself, cannot aspire at providing a respectable syntactic claim. This conclusion has immediate effects on the value of using an event-analysis in order to capture bound readings. The event-analysis seems to offer nothing more than a *perspicuous truth-conditional representation* of what is going on in cases of binding, but it can offer no *syntactic analysis* in the technical sense.

This should not be surprising: After all, this is what event-analyses are meant to do in general. As Terence Parsons suggests in the following passage, event-analyses have a clarificatory rather than empirical scope:

> I seek a theory that describes the semantics of sentences of English, that is, the relations between words of our language and things in the world. One convenient way of accomplishing this is to find a way to associate "logical forms" with English sentences. These "forms" will be sentences of a formal language that has already received a clear semantic treatment. The semantics of the English sentences in question will then be that of their associated logical forms. This intermediate route from the English to its semantics—via logical forms—is for the sake of convenience only. The semantics of the formal language I use (mostly the ordinary predicate calculus) is already widely known, and so I can presuppose a great deal of familiar work in formal semantics. (Parsons 1990: 11).

It appears evident to me that Parsons is here using a "revisionary" notion of logical form, as Stanley would call it. Sentences in the predicate calculus are used as "proxies" for the display of the truth-conditions of sentences in the natural language; their role is that of making the truth-conditions of these natural language sentences as accessible as possible to the theorist. As helpful as this can be, it seems evident to me that it doesn't and it cannot lead to any substantial syntactic, empirical claim. This implies that, even with relation to binding, the event-analysis responds to a need for clarification of the phenomenon, but it doesn't ultimately support any claim as to how things actually stand with respect to the real underlying syntax of the language. The conclusion is that, just like the Indexicalist's proposal, the Davidsonian proposal in terms of event-semantics does not and cannot probably aspire to be a genuine syntactic theory.

These considerations may trigger the following question: Is an explanation in terms of event-semantics then what Stanley addresses as "the trivial claim that, for purposes of interpretation, one should replace natural language by a notation in which all context-dependence is made explicit in the favored notation."? Have proponents of the event-analysis made just a minor discovery, with no empirical value? I do not think so. I think that one can regard the accomplishments of event-analysis as at least revealing something at the empirical level, namely with respect to the *conceptual resources* that are employed in the interpretation of sentences like "Everywhere I go, it rains".

My insight is that an event-analysis, even though it doesn't straightforwardly open the way to any syntactic claim, tells something of empirical relevance about how we interpret predicates and sentences by making explicit the way we organise the *concepts* that are invoked by any use of these linguistic items. Every time we use the predicate "rain", the correct use of the concept RAIN that we thereby invoke is such that it requires that every event of raining is located at a certain time and in a certain location; that is, the concept of RAIN is associated with *conceptual constraints* concerning which restrictions have to be observed in order for the use of the concept to count as correct. One can therefore view spatio-temporal aspects as already "packed into" the notion of event that underlies our talk of rainings. These aspects do not determine "argument-places" or "slots" in the

logical form of the predicate "to rain"; they do not reveal any syntactic structure; they are constitutive of the concept of RAIN, and operate at a purely conceptual level. Nothing, however, stands in the way of associating these conceptual constraints with linguistic representations in a formal language like that of predicate calculus. This can be done for the sake of convenience, clarity and perspicuity. One can therefore cash out talk of rain as involving quantification over event-variables, therefore obtaining sentences like (26). Once one introduces event-variables, it becomes apparent that this method of representation allows also to make sense of bound readings, as in (30), (31), (33). The tools of event-semantics allow the theorist to regard expressions characterising events as interacting with expressions that quantify over times and locations. The only assumption needed is that the quantifier expressions bind time-variables t and location-variables l that determine functions from these ts and ls to events occurring at (or relevantly related to) these ts and ls. This assumption is exactly what is presupposed by Cappelen & Hawthorne's treatment of binding in the event-analysis framework (see **section 10**).

An event-analysis should not, therefore be dismissed as "trivial", because devoid of any empirical relevance; in fact, it can be treated as an empirical hypothesis, even though not one concerning syntax. Event-analyses suggest empirical hypotheses concerning *the way we employ concepts*, and they do that by shedding light on conceptual relations, as for example relations between a rain event with its temporal/spatial aspects and other events with other temporal/spatial aspects, as in "Everywhere I go, it rains". An explanation of binding within the event-analysis may therefore commit the theorist to substantial empirical hypotheses, even though not in the same area in which the Indexicalist is working. The event-analysis still qualifies as a relevant alternative to Indexicalism if one accepts the methodological points made by Neale *at alia*, and agrees that no respectable *syntactic* theory is forthcoming from purely interpretive considerations like those concerning binding. The relevance of an event-analysis lies in its potentially illuminating aspects of our conceptual (and, for competent language users, also *semantic*) competence, i.e. aspects of our capacity to organise thought (and discourse) about certain matters.

Conclusion

In this chapter, I focussed on the issue whether conceptual constraints could be regarded as hidden positions in logical form, on the model provided by Indexical Contextualism. By arguing against Indexical Contextualism itself, my aim was to indirectly show that this account is not a good way of cashing out conceptual constraints. Throughout the chapter I proceeded in the following way:

- I introduced two examples of Indexical Contextualism, one due to Szabó, the other due to Rothschild and Segal. My view was that the differences between these two versions of the theory are marginal, to such an extent that they can be considered as mere variants of the same formal tool;
- I then moved to the evidence that has been offered as a ground for Indexicalism, which consists of a series of examples involving the phenomenon of binding, i.e. of co-variation of the values of a variable with the values introduced by a quantifier-operator. The conclusiveness of the evidence hangs on the assumption (called "Binding Assumption") that every case of semantic binding coincides with a case of *syntactic* binding;
- I surveyed four objections that have been levelled against the argument from binding. The first objection points at the inconclusiveness of the binding strategy: the argument from binding fails to establish the strong conclusion that *un-embedded* occurrences of "It's raining" carry a hidden argument position for rain-location, and only establishes—if it establishes anything—the weaker and less exciting conclusion to the effect that *embedded* occurrences of "It's raining" contain such a hidden position;
- The second objection charges the binding strategy of overgeneration: since bound readings can be reproduced indefinitely, this entails (*via* the Binding Assumption) that an indefinite number of hidden indexical positions can be posited. This entails sanctioning new

dimensions of context-dependence for the expressions at issue, which may go against our semantic intuitions. Overgeneration challenges are threatening for the Indexicalist only insofar as they take as their point of departure genuine cases of semantic binding;

- The third objection has to do with the fact that the alleged hidden indexicals do not share some of the basic features of genuine, non-hidden indexicals, such as perspectivality, giving rise to a priori truths and taking part to anaphoric relations;

- The fourth objection is a methodological point: the binding argument reaches conclusions about syntax by making considerations about semantics. This is not what orthodox methodology in syntax prescribes. As a result, Indexicalism cannot be approved as a fully respectable syntactic theory;

- My argument against Indexicalism starts from consideration of the phenomenon of *zero-readings*, i.e. readings of e.g. "It's raining", "Jill is ready", "The leaves are green" in which the predicates "to rain", "being ready" and "being green" are felicitously used as monadic;

- I argue that Indexicalism cannot capture zero-readings. Given that, according to the Indexicalist, the predicates mentioned above all include a hidden argument-place in logical form, it follows that sentences that give rise to zero-readings are all *open sentences*, i.e. sentences that are unevaluable, absent an assignment for variables. But sentences associated with zero-readings are not unevaluable—they are perfectly intelligible and evaluable. The Indexicalist can explain this fact only by supposing that either the variables receive a "free" contextual assignment of values or that they get bound by a quantifier. The first explanation is fully inadequate as an account of zero-readings; the second seems fine, but it is not available exclusively to the Indexicalist. Even an non-Indexicalist could have the resources to work out the bound reading. The real problem for the Indexicalist is that he cannot rule out the other reading—the one in terms of open sentences, something which can be successfully done by the theorist who envisages predicates like "to rain", "being ready" etc. as exhibiting a *monadic* syntactic structure. Ultimately then, the non-Indexicalist can explain the fact that zero-readings seem to involve monadic predicates with the idea that they *do* involve monadic predicates; this explanation is, quite obviously, not available to the Indexicalist (who assigns a more complex structure to these predicates).

- I illustrate how a Davidsonisn event-semantics is exactly what we need in order to account for zero-readings. Event-analyses also can accommodate bound readings, as Cappelen and Lepore (2007) show;

- It could be objected that, if Indexicalism cannot be taken seriously as a syntactic option, neither can the event-analysis. Proponents of Davidsonian accounts should therefore be content with having provided a mere logico-philosophical analysis of binding. I argue that this is not true. Even if event-analyses do not allow any syntactic claim, they may still commit their proponents to claims regarding the structure of our *conceptual (semantic) competence*. Adopting an event-analysis is not simply a way of achieving clarity through the use of formal notation; it can be endorsed as a way of making empirical claims as to how we organise concepts in thought and discourse.

Chapter 4

Conceptual Constraints and Relative Truth

Let us take stock. In **Chapter 1,** I defended the view that semantic under-determinacy is distinguished from other phenomena of semantic under-specification by the fact that the former, unlike the latter, is explainable in terms of just under-articulation. After having faced this foundational issue, in **Chapter 2** I moved to the problem of how communication is possible, despite the semantic under-determinacy of "It's raining", "Jill is ready", "The leaves are green", etc. There I introduced the idea that comprehension of utterances of such sentences is guided by conceptual/semantic constraints rather than by purely pragmatic demands. In **Chapter 3**, I focussed on the possibility of regarding conceptual constraints as hidden indexical elements, therefore examining a view known as Indexical Contextualism. My arguments focussed mainly on the motivations that have been offered to support an Indexical Contextualist solution, and in particular on whether Stanley's "Binding Argument" is genuinely compelling. I went through a number of reasons that theorists offered to emphasise that Stanley's argument is not entirely successful, and raised two points of my own. The first point has to do with the fact that Stanley's argument from binding seems to be driven by considerations of *uniformity* whose application is contentious, and which give rise to the unwarranted background assumption to the effect that, if a certain syntactic pattern is to be postulated in sentence S as embedded within operator Q, then that same syntactic pattern has to be postulated in sentence S when unembedded. The second point relied on the acknowledgment of so-called zero readings, as e.g. readings in which "It's raining" is regarded as a monadic predicate. I argued that the approach defended by Stanley would issue wrong predictions as to the truth-conditions of zero-readings. I subsequently adopted an *event analysis* of zero readings. Incidentally, by adopting this analysis theorists have been able to re-frame binding phenomena in a way that doesn't lead to Indexical Contextualism. I have argued that adoption of an event analysis doesn't force one to any syntactic commitment as to convert syntactic positions, even though it still has an empirical import: Event analyses may prove fruitful in order to represent the *structure of conceptual/semantic constraints*.

However, discarding Indexical Contextualism does not establish a conceptual constraints approach in a conclusive way. The theorist who is keen on reducing conceptual constraints to already acknowledged components of a formal semantic system may flag the possibility that conceptual constraints may receive yet another formal representation. On the account I will scrutinise in this chapter, conceptual constraints are to be represented as parameters to which the *truth-value* of under-articulated sentences is made relative. So for example, the conceptual constraint according to which rain has to fall somewhere is to be represented as a locational parameter *l* in the sentence's circumstances of evaluation. The constraint operates in such a way that, for every utterance of "It's raining", the *truth-value* (*and not* the content) of such an utterance is going to be relative to the value taken by the *l*-parameter in the relevant circumstances of evaluation.

Accounts in terms of relative truth for more or less restricted classes of sentences have been already developed in several areas of research in semantics and philosophy, mostly in very recent years. In what follows, I'll consider the different components that make up a proposal that one may want to call "Relativist", and I shall test their adequacy to provide an account of the role of conceptual constraints in utterance comprehension.

1. The Basic Move: Extra Parameters

In this section, I shall introduce the "basic move" that characterises a semantics typically called Relativist. As we will see throughout the chapter, this "basic move" characterises not only Relativism in its most proper formulation; it also underlies a form of Contextualism that is usually called "Non-Indexical". It is important to present the basic move, Relativism and Non-Indexical Contextualism separately, so as to avoid any conflation between the theories and the formal framework that lies behind them. Though they share the basic move, Relativism and Indexical Contextualism have significantly different implications, and should not therefore be taken as two mere variants of the same theoretical configuration.

The background apparatus for the "basic move" is the picture of semantics adopted by theorists since the work of Kaplan (1989), Lewis (1980) and Stalnaker (1970), which can be thought of as providing a *two-tiered* semantic machinery. At a first stage, the semantics takes as its input a sentence s together with the context c in which it is uttered: the input may be represented as a couple $<s, c>$. The $<s, c>$ couple is mapped onto the content that s expresses in c. Content is usually formally cashed out as an *intension*[30], i.e. a function from circumstances of evaluation to extensions. At the second stage, the intension maps the utterance's circumstances of evaluation onto an *extension*, i.e. the the truth-value that the utterance has in those circumstances.

It is important to distinguish, within this framework, the notions of *context of utterance* and of *circumstance of evaluation*. The role of a *context of utterance* is that of determining the content of all those expressions which do not have a fixed content across utterances. These are the so-called *indexical* expressions, whose content typically varies according to how a certain *feature* of the context of utterance varies. Thus, the content of "here" varies with the place of utterance, the content of "now" varies with the time of utterance, the content of "I" varies with the utterer and so on. A context could be represented as an n-tuple of coordinates $<u, t, p>$ each standing for one feature: u for the utterer, t for the time, p for the place and so on. On the other hand, a *circumstance of evaluation* has the role of determining a situation in which the *truth-value* of the utterance can be fixed. A circumstance of evaluation will consist of a set of parameters, though the question may arise as to what parameters may be included in a circumstance of evaluation. In general, whether a parameter is part of the circumstances of evaluation depends on whether there is an intensional operator which can affect evaluation. Let me illustrate. Suppose that an expression such as "Possibly" embeds a sentence such as "Pigs fly". This expression behaves as a sentential operator: it takes a sentence s and forms a more complex sentence of the form "Possibly, s". More importantly, this is an *intensional* operator insofar as it takes sentence s and "looks" at the truth-value of s at a possible world w_i (accessible from the actual world @). In other words, the operator determines a function from circumstances of evaluation in which the world-parameter w has been shifted to a possible world w_i (accessible from the actual world @), to truth values. An analogous reasoning applies to temporal and locative expressions such as "in the past", "somewhere", and so on. These expressions all take a sentence s and "look" at the truth-value of s at some time t_i (preceding the time of utterance t_u) or at some location l_i (different from or identical with the location of utterance l_u). In doing this, they determine functions from circumstances of evaluation whose parameters may have been "shifted" to a time t_i or a location l_i, to extensions.

Which formal representation should be given of a circumstance of evaluation? A circumstance of evaluation is standardly represented as an n-tuple of coordinates, and it usually includes only a possible world-coordinate, the corresponding index being written as $<w>$; more parameters may be added, for example temporal parameters t and locational parameters l, but this depends on one's

---

[30] It is not clear whether the content of a sentence s in a context c is also a proposition. Kaplan (1989) believes that contents are propositions in the sense of "what is said" in a context of utterance. Stalnaker (1970) also believes that contents are propositions in the sense of being the objects of our attitudes. By contrast, Lewis (1980) argues that contents cannot be propositions (what is said, the objects of our attitudes) because propositions do not behave compositionally. That's the reason why Lewis' semantics, though it envisages the level of context, it doesn't envisage a level of the content-in-context.

view concerning the corresponding intensional operators.

Monadic truth for an utterance of a declarative sentence *s* in a context *c* can be spelled out with the help of the notion of *truth at the index of the context*, in this way: a sentence *s* uttered in a context *c* is true *tout court* iff it is true at the index of the circumstances of evaluation *i* of the context *c*. If the index of the circumstances of evaluation includes just a possible world parameter *w*, a sentence *s* in a context *c* will be true *tout court* iff it is true at the world of the context w_c.

The fundamental traits of the underlying semantics having been outlined, the "basic" move that Relativists and Non-Indexical Contextualists share is quite easy to pin down: it consists in *positing extra parameters in the circumstances of evaluation of the sentence at interest—in addition to the possible-world parameter w—with respect to which truth-value is made to vary*.

An example of *ante litteram* relativist approach is *Temporalism* (see Prior 1968, Kaplan 1989). According to Temporalism, the truth-value of a sentence like "Carl is showering" is relative not only to a possible world *w*, but also to a time *t*. This implies that the proposition *that Carl is showering* is invariant in content across times, though it can be true or false depending on the value of the time-parameter in the circumstances of evaluation. For example, the proposition may be true at time t_1 iff Carl is showering at $<w, t_1>$ and false at t_2 iff Carl is not showering at $<w, t_2>$.

The model exemplified by Temporalism has been replicated in giving the semantics of other expressions in various areas of discourse. In general, the standard move is that of adding an extra coordinate in the circumstances of evaluation, corresponding to the feature of context truth is supposed to be relative to. Thus for example, Lasersohn (2005) argues that the truth of utterances of sentences containing predicates of personal taste like "tasty" and "delicious" is to be relativised to a world-time-individual triple $<w, t, i>$ rather than to a world-time couple $<w, t>$. This, according to the author, is a fully legitimate modification of the Kaplanian system:

> Kaplan treated the contents of sentences as (characteristic functions of) sets of time-world pairs. Contexts were assumed to provide a time and world, and a sentence *N* was defined as true in a context *c* iff the time and world of *c* were in the content of *N*. [...] What I would like to suggest is simply this: Instead of treating the content of a sentence as a set of time-world pairs, we should treat it as a set of time-world-individual triples. We assume that the context will provide an individual to be used in evaluating the sentences for truth and falsity, just as it provides a time and world; hence a sentence may be true relative to John but false relative to Mary. (Lasersohn 2005: 662-3)

An extra-parameter approach has been adopted to deal with epistemic modals, too. Egan, Hawthorne and Weatherson (2005) argue for the postulation of truth relative to *centered worlds*, i.e. to triples $<w, t, i>$ of a world, a time and an individual in order to account for the following phenomenon related with the use of epistemic "might" (and "must). Suppose that Moriarty asserts "Holmes might be in Paris"; Watson is in London, listening to Moriarty while unseen. Watson knows Holmes is not in Paris. He can concede that Moriarty spoke truly with his utterance, and yet this would not entail for him that, since Holmes might be in Paris, he (Watson) doesn't know that Holmes is not in Paris. An extra-coordinate approach to the truth of epistemically modal statements allows the theorist to say that, since the truth-predicate is relative to $<w, t, i>$, i.e. the ascription of truth to a certain assertion is always performed with respect to an centered world, from the fact that Watson acknowledges that Moriarty spoke truly, it doesn't follow that he spoke truly with respect to a Watson-centered world; so the acknowledgement of the appropriateness of Moriarty's assertion doesn't threaten Watson's knowledge of Holmes' whereabouts, because such acknowledgement is, so to speak, relative only to a Moriarty-centered world. Similar treatments of epistemic modals, based essentially on the addition of one or more parameters in the index of the circumstances of evaluation, can be found in Egan (2007) and von Fintel and Gillies (2008).

Kompa (2002) gestures at a treatment in terms of extra-parameters for epistemic verbs like "knows": as she argues, "know" is to be regarded as context-sensitive, though not as indexical in the same way as "here" (contrary to what Epistemic Contextualists such as Unger (1984, 1986), DeRose (1992, 1995) and Cohen (1986, 1987) claim). What changes from context to context is not the content of the word "know"; i.e., it is not the case that a different knowledge-relation is expressed in each different context. Rather, the interests, purposes and goals of the speaker determine a standard of knowledge which affects the *truth-value* of knowledge ascriptions. Kompa's informal considerations receive a formal organization by MacFarlane (2005b). The first step in MacFarlane's strategy is that of positing an extra epistemic parameter e in the circumstances of evaluation, with respect to which the truth of assertions of the form "A knows that P" varies. Only after having done the "basic move" he claims that a knowledge attribution "A knows that P" is true iff A knows that P at the world w_u of the context of utterance and at the epistemic standard e_a of a special context, which he calls "context of assessment", i.e. a context in which the proposition as uttered at w_u is evaluated.

The extra-parameters move has been adopted in order to cope with semantically under-determined expressions as well, such as "green" or "ready". At the root of the strategy lies a general idea (defended e.g. by Predelli 2005a, b) about the relation between semantically under-determined expressions and context, a very straightforward development of which is certainly the postulation of extra-parameters. On the one hand, traditional semantic systems such as those designed by Kaplan and Lewis are concerned with assigning functions from "points of evaluation" to truth-values to pairs <s, c> consisting of a clause s (or interpreted logical form) and an index of the context c (a list of parameters: agent, time, location). Provided absence of indexicality, different utterances of the same sentence certainly have the same intension. Nevertheless, variability of truth-value despite sameness of intension seems possible. In Travis' notorious green leaves example, the sentence "The leaves are green" is uttered by Pia in two different occasions within the same context: on the first occasion, while talking with with a photographer, in the second occasion, while talking with a botanist. The two utterances have intuitively different truth-values, though obviously being assigned the same intension by the semantics. What could cause the different evaluations? According to Predelli, the conversants' interests and standards determine different "applications" of the intension, that is to say that they determine different sets of circumstances in which the utterance might be true/false (see 2005a: 366). This implies that Pia's utterance is true in the first occasion because the photography-related interests that are operative in the conversation determine circumstances which are such as to make the sentence "The leaves are green" true. Pia's utterance is however false in the second occasion because, given the botanist's interests and standards, the circumstances determined by such interests are such as to make the sentence "The leaves are green" false. Predelli's idea is that of maintaining the traditionally accepted Kaplan-style semantics while allowing that the same proposition be evaluated at different circumstances, whose configuration and composition is shaped by the conversants' interests and standards. As one may realise, the idea that the conversationalists' interests, goals and presuppositions determine different circumstances is altogether neutral with respect to the structure one may want give to the index of the circumstances at the formal level. Predelli remains uncommittal in this sense: he is happy to still regard these circumstances as "possible worlds", if by "possible world" one merely means a "point of evaluation". Instead of using possible worlds, one may as well use world-time pairs, or even situations in the sense of Barwise and Perry (1983): for example, authors such as Corazza & Dokic (2007, 2010) Barba (2007) and, to a certain extent, Gauker (2010) take this latter option. The technical details are not important as the general underlying idea. What matters to us is that the semantics of expressions such as "green" or "ready" might be formulated by tinkering with the parameters included in the circumstances of evaluation of a Kaplan-Lewis-Stalnaker machinery—something which appears feasible at first sight.

Indeed, the path is short between devising a solution such as that set forth in broad outline by

Predelli and adopting extra-parameters in the circumstances of evaluation. MacFarlane (2007a, 2009) suggests such a strategy, therefore explicitly opting for the "basic move", i.e. therefore explicitly sanctioning a strategy in terms of extra-parameters in the semantic treatment of semantically under-determined expressions. As he writes, in order to account for the truth conditions of e.g. "The leaves are green" or "Chiara is tall", we can change the way we think of circumstances of evaluation.

> Possible worlds will presumably be one component of our circumstances of evaluation (otherwise, what will our modal operators shift?) but nothing stops us from introducing other components as well. […] So let's think of a circumstance of evaluation as an ordered pair consisting of a world and a "counts-as" parameter, which we can model as a function from properties to extensions (function from worlds to extensions). The "counts" as parameter is so called because it fixes what things have to be like in order to count as having the property of tallness (or any other property) at a circumstance of evaluation. (MacFarlane 2007a: 246)

MacFarlane's way of cashing out the circumstances of evaluation for utterances of "The leaves are green", "Chiara is tall", etc. is thus in line with the "basic move", insofar as it regards the truth-value of these sentences in a context as being relative not only to a possible world w, but also to a "counts-as" parameter c, i.e. to a standard which determines, for each occasion, what falls and does not fall in the extension of a relevant predicate ("green", "tall"), thus affecting the overall utterance's evaluation.

Thus far, we have simply reviewed the various cases in which the basic move has been opted for, without mentioning whether this is or not problematic. Though the authors that subscribe to it are confident that the move is not particularly problematic (*prima facie*, it doesn't lead to any contradictory nor counterintuitive results), some have worried that it's not sufficiently motivated. The reason is that, as I mentioned previously, the addition of extra parameters has been justified by its original proponents (e.g. Kaplan and Lewis) with the presence of *operators* which are capable of shifting the value of these parameters. As Kaplan and Lewis explicitly write:

> A circumstance will usually include a possible state or history of the world, a time, and perhaps other features as well. The amount of information we require from a circumstance is linked to the degree of specificity of contents, and thus to the kinds of operators in the language. (Kaplan 1989: 502)

> Often the truth (-in-English) of a sentence in a context depends on the truth of some related sentence when some feature of the original context is shifted. "There have been dogs" is true now iff "There are dogs" is true at some time before now. "Somewhere the sun is shining" is true here iff "The sun is shining" is true somewhere. "Aunts must be women" is true at our world iff "Aunts are women" is true at all worlds. "Strictly speaking, France is not hexagonal" is true even under low standards of precision iff "France is not hexagonal" is true under stricter standards. [...] I emphasized that context-dependence was multifarious, but perhaps the shifty kind of context-dependence is less so. The list of shiftable features of context may be quite short. I have suggested that the list should include time, place, world, and (some aspects of) standards of precision. I am not sure what more should be added. (Lewis 1980: 27)

As Lewis and Kaplan grant in the above passages, the justification for adding a world-coordinate in the circumstances of evaluation is provided by the presence of what, in formal terms, is called

modal operators, like "possibly" and "necessarily" (but also the modal "must" and "might"). These operators are described as "shifting" the world of evaluation. This allows for the assessment of sentences like "Possibly *p*" and "Necessarily *p*". The semantics of these sentences standardly requires that one "checks" whether *p* holds at worlds $w_1, w_2, \ldots w_n$ accessible to @ (where @ could belong to this set too). Since there is a strong *prima facie* analogy between modality and tense, Temporalism has been built up on analogous grounds. The addition of a time-coordinate in the circumstances of evaluation has been justified by the presence of *temporal operators*, like "It was the case that" and "It will be the case that". These operators are described as "shifting" times, thus making it possible to evaluate sentences such as "It will be the case that *p*" and "It was the case that *p*". Roughly, the assessment of such sentences requires that one "shifts" the time of evaluation from t_n to t_{n+m}, in the case of future operators, or from t_n to t_{n-m}, in the case of past operators, and checks whether *p* holds at those times. Similar considerations hold for locational parameters, which are allegedly justified by operators such as "somewhere", and for standards of precision parameters, which are allegedly justified by the presence of expressions like "strictly speaking".

Now a question may arise whether the presence of operators is the only way of justifying the addition of extra-parameters. Are the cases of modality and tense paradigmatic? If they are, then anyone who wants to accomplish the basic extra-parameter move in an "orthodox" way should be able to justify the addition of an extra-parameter by displaying some operator in the language capable of shifting the value of that parameter. The problem is that not all the recent proponents of the basic move can do that, or at least it is questionable that they can do that. For example, with respect to undertaking the basic move on "know", Stanley (2005a) points out that, if we wish to countenance an epistemic standards parameter in the circumstances of evaluation, there must be an operator in our language that is purported to shifting it; but in the case of knowledge predicates, there seems to be no such operator[31]; it would be odd for the Relativist to maintain that knowledge ascriptions are true relative to an epistemic parameter while at the same time acknowledging that there are no epistemic operators; for this reason it would be wrong for the Relativist to postulate such an epistemic parameter in the first place. The postulation of extra-parameters in this area of discourse would turn out as blatantly unjustified and hence not completely respectable as a semantic theory of "know". One can see that the structure of the argument is general enough in order for it to be extended to other versions of the basic move—e.g. the introduction of an epistemic parameter *e* for epistemic modals, of an individual parameter *i*, or a counts-as-parameter *c*. I shall leave this task to the reader.

It is not my interest at this point of the chapter to assess whether the various versions of the basic move are grounded enough on considerations as to the presence of operators that may "shift" the features that these parameters represent. I will just assume, for the sake of my arguments against Relativism and Non-Indexical Contextualism, that the basic move is motivated or at least there are prospects for motivating it. In the rest of this chapter, my attention will mainly be devoted to assessing, on the basis of other kinds of arguments, whether conceptual constraints can be traced to

---

[31] The Relativist who wishes to point at such epistemic operators may be tempted to refer to the work of Ludlow (2005). In the paper, Ludlow defends a contextualist position about the semantics of knowledge ascriptions by noting that there are felicitous utterances of, for instance, "She doesn't know that *p*, *by the standards of chemistry*". Here the expression "by the standards of chemistry" is supposed to make the epistemic standards for "know" explicit. If this is so, then this suggests that there's room for such epistemic standards at the level of the truth-conditions of knowledge ascriptions, even when speakers do not explicitly articulate them. The Relativist could take advantage of Ludlow's results and claim that expressions such as "by the standards of chemistry", "according to John" etc. ought to be regarded as epistemic operators. Stanley (2005a: 164) opposes this move (whether or not advanced by the Relativist), and argues that these expressions do not have any special epistemic import, though they can be felicitously appended to knowledge ascriptions, thus forming sentences like "Mary knows that *p*, *by the standards of chemistry*", or "*According to John*, Mary knows that *p*". This is easily demonstrated by the fact that the same expressions can be appended to sentences that do not contain any epistemic terminology, and therefore cannot be used to make knowledge ascriptions. For example: "*By the standards of chemistry*, what is in the Hudson river is not water"; "*According to John*, the bus will arrive soon".

anything similar to the extra-parameters introduced *via* the basic move.

2. Extra Parameters and Conceptual Constraints

I take conceptual constraints to be ways of structuring and organising thought and discourse on determinate matters; as I argued in **Chapter 2**, there is reason to believe that these conceptual (semantic) constraints guide us in our understanding of utterances of semantically under-determined expressions. Since the notion of conceptual constraint does not presuppose any particular formal framework, towards the end of the second chapter I speculated that there could be more than one way to represent conceptual constraints within different formal semantic theories. After having explored and discarded the idea that conceptual constraints may be hidden indexical elements (see **Chapter 3**), the idea that we are currently exploring is whether conceptual constraints can be seen as parameters of evaluation. In this section, I will be mainly concerned with providing examples of how this particular implementation of conceptual constraints can be accomplished.

Let us consider an occurrence of "It's raining". Conceptual constraints on the concept of RAIN dictate that every instance of a sentence that contains the predicate "to rain" be true iff the raining is occurring at a location. This truth-conditional result ensues, I argue, from one's exercise of one's conceptual/semantic competence on the concept RAIN (and the predicate "to rain"), which is such that it dictates that the standard format for thought and discourse about matters of rain necessarily contemplates a locational aspect. Since, in the framework at issue, constraints have to operate on the truth-value of sentences, this entails that one's application of conceptual constraints makes it that the truth-value of the sentence "It's raining" will be relative to a locational parameter l, that is, any occurrence of the sentence "It's raining" will have the following truth-conditions:

(1) "It's raining" is true $<w, l>$ iff it's raining at w, l.

Let us now consider the sentence "Jill is ready". In light of the conceptual/semantic constraints that govern the use of the concept BEING READY (and the predicate "being ready"), any competent user of this sentence may be in a position to determine that instances of sentences like "X is ready" are true iff X is ready for a certain purpose, or activity. This truth-conditional result is putatively obtained by means of one's exercise of one's competences on the readiness concept/predicate, insofar as the standard cognitive/discursive format on the notion of being ready requires necessarily a purpose- or activity- element. Since in the framework currently explored, conceptual constraints must correspond to parameters of evaluation, this entails that, by applying one's conceptual/semantic competences, one will be able to determine that any occurrence of the sentence "Jill is ready" will have its truth-value relativised to a purpose-parameter p. The related truth-conditions will be expressed as follows:

(2) "Jill is ready" is true $<w, p>$ iff Jill is ready at w, p.

Let us now consider "The leaves are green". Correct uses of the concept of BEING GREEN (and the related predicate "being green") should comply to the constraint that greenness is attributed to an object X under a certain respect. Application of such a constraint by a competent concept/word-user is expected to imply that every instance of "X is green" will be true iff X is green in a certain respect. In the framework under consideration, the constraint has to operate on truth-value; this entails that instances of "X is green" will have a truth value relative to a respect r, in the following way:

(3) "The leaves are green" is true $<w, r>$ iff the leaves are green w, r.

We have made some examples of how an extra-parameters strategy could be used to implement the

idea that conceptual (semantic) constraints seem to guide comprehension of semantically under-determined sentences. Is this strategy also advantageous? The benefits of seeing conceptual (semantic) constraints as parameters on the determination of truth-value may be more than one.

First of all, the extra-parameters move as a way of representing conceptual constraints doesn't commit the semanticist to any thesis concerning additional covert elements in the sentence's syntax. As we saw in **Chapter 3**, Indexical Contextualism cannot but be committed to the existence of positions marked at the level of deep syntax or logical form. The main reproaches directed at the Indexical Contextualist were that the positing of such positions in logical form is not completely motivated by the lights of the syntactician's methodology. As opposed to Indexical Contextualism, an account based on extra parameters is not threatened by this kind of criticism, for the claim that truth is relative to extra-parameters only affects the structure of the T-sentences associated with sentences of the natural language: its commitments are therefore not at a syntactic level, but at a *metalinguistic* level. All one has to commit to is that we switch from T-sentences in which, like in (4), truth is relative to just w, to T-sentences in which, like in (5), truth is relative to more parameters than just w:

(4) "It's raining" is true at w iff it's raining at w.
(5) "It's raining" is true at w and l iff it's raining at w and l.

This change may not be free of costs as well; however for the present purposes it's enough to note that a revision of the meta-language in which the truth-conditions of semantically under-determined sentences are expressed doesn't bring with it any unwarranted commitment to extra syntactic structures at the level of the object language.

Second of all, the extra parameters move is compatible with a semantic picture in which intuitive comprehension of an utterance of a semantically under-determined sentence is possible, even though there is no particular context from which to extract the elements that are missing from the sentence's truth-conditions. This is what proved impossible in a framework like that of Radical Contextualism (see **Chapter 2**), in which, unless a particular context is available, intuitive comprehension of a semantically under-determined expression was predicted as impossible. It is worth noting that, in an extra-parameters framework, the intuitive comprehension of a semantically under-determined sentence just depends on the speaker's mastery of (a) a compositional semantics for English; and (b) the conceptual/semantic apparatus associated with such a language. What the speaker needs to do is (a') associate the correct linguistic meaning to occurrences of the sentence and (b') master the related conceptual/semantic aspects enough in order to know that the truth of a sentence containing these-and-these terms is relative to these-and-these parameters. No knowledge of any particular context is required in order for the speaker to intuitively grasp the truth-conditions of an utterance of a semantically under-determined sentence. So for example, if the proponent of the extra-parameters move is right, in order for the speaker to intuitively grasp the truth-conditions of "It's raining", it would be sufficient that the speaker knows the content conventionally expressed by the sentence (<*that it's raining*>) and sufficiently masters the concept RAIN (and the related predicate) in order to know that any occurrence of "It's raining" is true or false only with respect to a circumstance which includes a location l. These two factors are sufficient in order to gain an intuitive knowledge of the truth-conditions of an utterance of "It's raining" even out of a context, i.e. knowledge of the conditions in which any *arbitrary utterance* of the sentence would be true or false. Analogous considerations apply for "being ready" and "being green".

In this section I have illustrated how what I call conceptual (semantic) constraints can be integrated into a Kaplan-Lewis-Stalnaker system by cashing them out as parameters in the index of the circumstances of evaluation. The view would have it that that conceptual (semantic) constraints determine the structure of the index of the circumstances of evaluation by providing a specific parameter. For instance, conceptual (semantic) constraints on "to rain" determine that an utterance

of "It's raining" is true at $<w, l>$ iff it's raining at w, l—where l is a location-parameter. The advantages of such a view are the following: (i) it avoids commitments to covert elements articulated at the level of syntax; (ii) it allows for intuitive comprehension (i.e. of comprehension of what the utterance says in any arbitrary context).

3. Non-Indexical Contextualism vs. Radical Relativism

As I already mentioned, what I have called the "basic move", consisting in the addition of an extra-parameter (or more extra-parameters) in the index of the circumstances of evaluation can be developed in two ways, one determining a position called Relativism, the other determining a position called Non-Indexical Contextualism. The two accounts share the extra-parameters insight, and they depart with respect to the question of *what fixes the value* of the postulated extra-parameter(s). Each account is suited to explaining different phenomena and to be contrasted with different rival accounts. In this section, I will present both Relativism and Non-Indexical Contextualism, by highlighting their differences in basic claims and implications.

The first account I shall present on is known as *Non-Indexical Contextualism* (but also as *Moderate Relativism*). Its central statements are:

[Extra-parameters]
The truth of certain sentences is relative to an extra parameter in the circumstances of evaluation;

[Context of Utterance]
The value of the extra parameter is systematically fixed *by the context of utterance*.

This theoretical position plays a role in several debates[32]. In order to get clear about the potentialities of the view, we shall therefore briefly survey the contributions that the position makes to the debates in which it is defended.

The notion of "Non-Indexical Contextualism" was elaborated originally (by MacFarlane 2007a, 2009) within a debate about the nature of the context-sensitivity of some expressions, as e.g. "red", "tall", "ready". The view was meant as an alternative to both Indexical Contextualism and sheer Invariantism about the semantics of these expressions, in that it would highlight a way for expressions to depend on context—as an Indexical Contextualist would want—by at the same time expressing the same content in every context—as an Invariantist would want. Non-Indexical Contextualism is a position that escapes the usual objections against Indexical Contextualism, and that explains the same range of data as Invariantism. Not only that: as a matter of fact, it is also the only position that arguments typically taken to support Indexical Contextualism really manage to support. These points are clearly illustrated by MacFarlane.

One typical argument that is levelled against Indexical Contextualism by Invariantists (as e.g. Cappelen and Lepore 2005, Hawthorne 2004) concerns the fact that, at least when it comes to expressions such as "tall", "ready", "know", we can make perfectly accurate homophonic reports of the utterances of other speakers, even when we are in a context that differs relevantly from the one of the original utterance as to one or more truth-conditionally salient elements. So for example, suppose that A utters "The cake is ready" in context C_1, meaning that it is ready to be eaten. B reports A's utterance in a context C_2 in which the cake's being ready to be thrown in someone's face

---

[32]   In the following sections, I shall only talk about two of the debates in which the position plays a role, even though actually there is a third field in which Non-Indexical Contextualism has been adopted, in the intersection between semantics proper and the philosophy of mind. So for example, Recanati (2007) adopts a form of Non-Indexical Contextualism (or Moderate Relativism, in his terminology) in order to capture some essential features of the content of *de se* thoughts, i.e. thoughts concerning the perspective of the self. The debate in which this position is defended is only marginally related to the question of semantic under-determinacy and context dependence, and I shall therefore not devote too much attention to it here (for basic reference see Lewis 1979, Perry 1977, 1979, Castañeda 1999, Evans 1982).

is discussed, by uttering "A said that the cake us ready". Still, even though the context has relevantly changed, the content of the reported sentence is not influenced: The report made by B when she utters "A said that the cake is ready" says something *true* about what A said. This means that the predicate "ready" doesn't exhibit any context sensitivity that one could equate with indexicality. Notice that this *cannot* happen whenever we are homophonically reporting a sentence which contains an indexical, like for example "I". In (6), the report made by B is clearly inaccurate in that, as context changes relevantly (in C_2 the speaker is B, not A), B's report comes to say something false concerning what A said:

(6) In C_1, A: "I'm ready"
    In C_2, B: "A said that I am ready"

As MacFarlane observes, the Invariantist's argument undercuts only the Indexical version of Contextualism, and it leaves the Non-Indexical variety untouched. As a matter of fact, Non-Indexical Contextualism can explain why B's report of "The cake is ready" is successful. The explanation is that, since it's not the content of the sentence "The cake is ready" that varies with context, but it's rather its truth-value, a difference in salient elements in the context of the report cannot have any effect on the sentence's content. The sentence "The cake is ready" has therefore the same content both in A's original utterance and in B's report: This is enough to explain the accuracy of B's report, i.e. the fact that B manages to say something true about what A said.

MacFarlane also observes that the main argument that has been put forward in order to support Indexical Contextualism in the end just supports Non-Indexical Contextualism. This argument is known as the Context-Shifting Argument (I myself covered it in **Chapter 1**). The structure of the argument is the following: one is invited to consider two utterances of the same sentence *s* in two different contexts C_1 and C_2. Intuitively, it seems that, in C_1, *s* says something true and, in C_2, *s* says something false. Proponents of Indexical Contextualism argue that if the two utterances differ in truth-value, they must differ in content too. Hence, the sentences must contain an indexical element. MacFarlane argues that this conclusion can be established only if the proponent of Indexical Contextualism can prove that the two utterances take place in the same circumstances. However, it is very likely that the circumstances of the two utterances differ, and the Non-Indexical Contextualist has the means to represent this change by construing the index of the utterance's circumstances of evaluation as containing an extra-parameter, which gets a different value in C_1 and C_2. In light of these considerations, and in absence of any demonstration from the part of the Indexical Contextualist that the circumstances of the two utterances do not differ, the only conclusion that follows from the Context-Shifting Argument is that the two utterances of *s* have different truth-values because truth-value depends on the different *circumstances of evaluation* of C_1 and C_2 respectively. Both sentences, however, express the same content across the two utterances, and therefore cannot contain any covert indexical element.

Non-Indexical Contextualism also plays a role in a parallel debate, focussed on the notion of *faultless disagreement* (Wright (2001, 2006), Kölbel (2002, 2004a, b)). (The fact that, in this debate, Non-Indexical Contextualism is often called "Relativism" should not mislead us, for the central claims of this view are still [Extra-parameters] conjoined with [Context of Utterance]). The debate on faultless disagreement revolves around the following problem: people often disagree on matters of taste, or on morality, or aesthetics; since taste, morality and aesthetics are amenable to subjective preferences or inclinations, disputes in this area have been called *disputes of inclination*. For example, the following are typical examples of disputes of inclinations:

(7) A: "Guacamole is tasty"
    B: "No, you are wrong. Guacamole is not tasty"

(8) A: "Eating meat is right"
     B: "I disagree. Eating meat is not right"

(9) A: "The Taj Mahal is beautiful"
     B: "Not at all. The Taj Mahal is not beautiful"

Now what is peculiar about disputes of inclination like those in (7)-(9) is that, on the one hand, they seem to involve genuine contradictions: after all, the speakers involved in the disputes use expressions like "No", "Not at all", "You are wrong", "I disagree". However, it also seems that none of the speakers needs to be making a mistake, i.e. it seems that neither of the speakers is at fault about how things are. The disagreements that take place in this area of discourse are therefore called *faultless disagreements*.

Words expressing properties depending on subjective inclinations ("tasty", "right", "beautiful") all give rise to examples of faultless disagreement. It is therefore a *desideratum* for any theory aiming at describing the semantics of these expressions that it captures both the element of contradiction and the element of faultlessness that characterise disagreement in this area of discourse, thus explaining the notion of faultless disagreement (or alternatively, explaining it away).

One option that, at first sight, captures well the faultless nature of utterances in disputes of inclination, is Indexical Contextualism (for a proposal that has been identified as going in this direction, see e.g Harman 1978). The Indexical Contextualist just needs to say something like this: predicates like "tasty", "right", "beautiful" are all associated with an extra argument-slot in logical form corresponding to a standard of evaluation (either a taste-standard, or a moral-code-standard, or an aesthetic-canon standard). This makes it that, for every utterance of e.g. "Guacamole is tasty", the content-in-context of the sentence is relativised to a standard of evaluation—perhaps that of the speaker, the resulting contents being, for instance, *that guacamole is tasty by the standards of A*, or *that guacamole is tasty by the standards of B*. Indexical Contextualism can therefore say that A and B are faultless, as long as it is true that guacamole is tasty by the standards of A as well as that guacamole is tasty by the standards of B. The problem is that Indexical Contextualism cannot say that the two speakers A and B are asserting contradictory contents, for the proposition that A's assertion expresses, namely *that guacamole is tasty by the standards of A* is entirely compatible with the proposition asserted by B, *that guacamole is tasty by the standards of B*. What is going on between A and B is very similar to what would obtain if A said "I wear glasses" and B denied it by saying "I don't wear glasses": even though they superficially contradict each other, at the level of the proposition expressed one can see that they are really just saying two different things (see Wright 2001, Kölbel 2004, Lasersohn 2005)[33].

In addition to this, Indexical Contextualism brings to a distorted conception of the topic of conversations in the areas of discourse covered by its treatment. Moving to the field of morality, for example, the subject matter of a discussion over whether it is right or wrong to eat meat would not be, according to Indexical Contextualism, the issue of eating meat, but it would rather be each speaker's moral code. However, we do not generally take speakers who are engaged in a discussion

---

[33] Indexical Contextualism can be amended so as to be able to account for the contradiction element. The trick consists in dropping the idea that the content of, e.g., "tasty" is fixed by the standards of taste of each individual speaker. Some theorists have suggested a version of Indexical Contextualism according to which the content of "tasty" varies relative to the standards of a collective entity, such as the conversationalists' group, or their community of reference (Glanzberg 2007, Sundell 2011). This looks like a viable proposal, even though doubts may be raised as to the conversational dynamics it gives rise to (see Lasersohn 2005, Sundell 2011). Alternatively, the Indexical Contextualist may preserve the idea that the content of each utterance of "tasty" is speaker-dependent; to rescue disagreement he has to posit that the parties involved in the dispute are *presupposing a common standard* (see Lopez de Sa 2008,). Concerning this proposal, it's not clear whether it accounts for real disagreement or merely for the appearance of disagreement which the conversationalists determine in their situation (see e.g. Coliva and Moruzzi, MS); moreover, it doesn't seem that the speakers involved in a taste dispute necessarily presuppose that they have the same taste-standards (see Baker 2012).

on morals to really be talking about their own moral codes: it seems that there's a common conversation topic they are focussing on (whether eating meat is or is not right) and they express their judgements about (see Kölbel 2004b).

Non-Indexical Contextualism seems to solve the difficulties that afflict Indexical Contextualism. First, it makes sense of the intuition that there is a contradiction between the contents asserted by A and B for, according to Non-Indexical Contextualism, the proposition that A asserts and B denies is the same (namely, *that guacamole is tasty*). A and B are not asserting two distinct contents; their exchange focuses on the very same proposition, *that Guacamole is tasty*, and contradiction arises to the extent that A accepts it while B accepts the negation of it. Secondly, Non-Indexical Contextualism doesn't distort conversation topics. Since the content-in-context of sentences is not specific as to any standard of evaluation, this entails that when people are discussing taste or morals, the topic of each conversant's assertion is not his or her own taste-standards or moral-code-standards, but it is rather the very issue being under discussion, e.g. the rightness of eating meat.

Despite the seeming superiority of Non-Indexical Contextualism over the Indexical account, some theorists have argued that the benefits offered by Non-Indexical Contextualism do not ultimately help capture the contradiction element that characterises faultless disagreement. MacFarlane illustrates this point in his MacFarlane (2007b: 25). His reasoning goes as follows: even though the truth of what A says (*that guacamole is tasty*) entails the falsity of what B says (*that guacamole is not tasty*) in any circumstance of evaluation, nevertheless Non-Indexical Contextualism states that the contents asserted by the two speakers are to be evaluated at the circumstances of each speaker's context of utterance. We may say that, if the speakers are both sincere, the content *<that guacamole is tasty>* is true at A's context (which contains A's taste-standards), while the content *<that guacamole is not tasty>* is true at B's context (which contains B's taste-standards). But then, since each content is true at each speaker's circumstance of evaluation, there is no contradiction, and hence no disagreement.

A similar point holds, according to MacFarlane, with respect to the phenomenon of retraction (see MacFarlane MS, chs. 1, 5). Performing a retraction means, informally, rejecting an assertion that one has made previously. For example, if at t_1 I believe that Carl is married with Paula, I may assert "Paula is Carl's wife". Suppose however that, at t_2, I learn that Carl is only engaged with Paula. Then I may be inclined to reject my previous assertion on the account that it was false to say that Paula is Carl's wife because Paula is just Carl's fiancée. Performing a retraction thus involves *reassessing one's previous claim*—for example, reassessing at t_2 my claim at t_1 to the effect that Paula is Carl's wife. Retraction poses a problem for Non-Indexical Contextualism, in the following sense: if, as Non-Indexical Contextualism has it, the truth of our utterances is fixed at the very context of our utterances, then it becomes impossible to re-assess the correctness of a previous act of assertion, for the correctness of that act is fixed once and for all at the context in which the assertion was originally made. Retraction would then be impossible, in that we could not re-assess our previous utterances: their truth-value would be unchangeable, because fixed once and for all by the circumstances of the context of utterance.

The distinctive claim of Non-Indexical Contextualism, to the effect that the truth of utterances is fixed at the circumstances of evaluation *of the context of utterance*, thus entails serious limitations in capturing phenomena such as disagreement in matters of inclination and retraction. These limitations are what motivated some authors to elaborate a refined version of Non-Indexical Contextualism, which for many deserves the name of Relativism. The second theory that stems from the extra-parameters approach that I will now describe is what is usually called just *Relativism* (or *Radical Relativism*). Its core claims are:

[Extra Parameters]
The truth of certain sentences is relative to an extra parameter in the circumstances of evaluation;

[Context of Assessment]
The value of the extra parameter is fixed by the so-called "context of assessment", which is to be kept apart from and is wholly independent of the context of utterance.

The evident novelty introduced (Radical) Relativism is the notion of "context of assessment". What is a context of assessment? Metaphysically speaking, a context of assessment is just like a context of utterance: it may be thought of as a set of parameters corresponding to various features, as for instance the individual who is doing the assessment, the place, time, taste-standards, etc[34]. The difference between a context of utterance and a context of assessment is in their role. In a context of utterance, a proposition is put forward through an act of speaking; in a context of assessment the assertion of a proposition is *simply assessed as true or false*. This distinction between context of utterance and context of assessment entails that one and the same speech act can be evaluated in an indefinite number of contexts of assessment. The context of assessment *determined by the context of utterance* will for sure have a central role in common assertive practices, however this role is not acting as a constraint on the correctness and incorrectness of assertions any longer: For in this framework it is legitimate to assess a certain assertion from a context of assessment that is completely unrelated to the context of utterance.

This gives good prospects for capturing both disagreement and retraction in the cases at interest. For suppose that A and B are disagreeing over whether guacamole is tasty. A asserts "Guacamole is tasty" and B asserts "Guacamole is not tasty". Now, the truth value of A's and B's assertions is not established by their contexts of utterance, but by a context of assessment. This entails that, from A's context of assessment, B's assertion is false iff A's is true, while from B's context of assessment, A's assertion is false iff B's is true. In other words, disagreement is rescued to the extent that, for any context of assessment c_a, A's and B's assertions cannot both be true (and therefore cannot be both accurately performed) at that context of assessment c_a.

The same holds for retraction. Suppose ten years ago I asserted "Guacamole is tasty". After ten years, my taste has changed and I now believe that guacamole is not tasty. Therefore, I may want to retract my assertion by saying that it is false, or incorrect. Given that the truth of my previous assertion is not fixed "once and for all" at the context c_u in which it was performed, but it is rather fixed at a context of assessment c_a, I can surely reconsider my previous assertion and re-assess it in the light of the new context of assessment I now occupy. This allows me to say that my former assertion is incorrect, rather than to say that, even though my former assertion was correct *then*, it would be incorrect *now*.

To sum up, in this section I have illustrated two different ways in which the extra-parameters move could be developed, called respectively Non-Indexical Contextualism and (Radical) Relativism. The two accounts share the common "basic move" consisting in the postulation of extra parameters in the circumstances of evaluation, though they depart in their respective views concerning what fixes the value of such parameters. Non-Indexical Contextualism has it that the value of any extra-parameters is fixed by the (index of the) *context of utterance*; on the other hand, Relativism has it that the value of the extra-parameters is fixed by the features of the so-called *context of assessment*. Each view is suited to accounting for different facts and has different implications. The purpose of this section was that of presenting each theory's most salient aspects and consequences in broad outline. In the next sections, I'll consider the idea of cashing out conceptual constraints as parameters in the index of the circumstances of evaluation—in the way of Relativism (**section 4**) and of Non-Indexical Contextualism (**section 5**) respectively.

---

[34] This is true only if we mean by "metaphysical similarity" similarity as to the "basic components" of which a context of assessment is made of (individuals, locations, times...). In this respect, a context of assessment and a context of utterance are metaphysically alike. This however may obscure the fact that contexts of assessment may consist of *combinations of features* that make them metaphysically very different from contexts of utterance. For example, a context of assessment may contain the actual world, but a time in which the speaker of the context of utterance doesn't exist at all (thanks to Marco Santambrogio for having made this apparent to me).

4. Conceptual Constraints and Relativism: a Motivated Move?

As mentioned in **section 3**, one of the main motivations that have been offered in favour of Relativism, for example by MacFarlane in many of his writings, is that this theory is particularly well-suited to capture the fact that two speakers involved in a dispute over what is tasty, beautiful or right are really disagreeing (despite the fact that each of them is in some sense issuing a claim which is correct on the basis of subjective considerations). Arguments in favour of Relativism also draw on the notion of retraction, whose structure is once again allegedly well captured by MacFarlane's theory. In this section, my aim is that of dwelling a little bit more on the Relativist's motivations. This is meant to connect back with the issue of conceptual constraints by means of the following question: are the Relativist's motivations relevant for the theorist who wishes to give an account of conceptual constraints in terms of relative truth? In other words, *supposing that the theorist aims at cashing out conceptual constraints in relativistic terms, would the move be motivated*?

Let me begin by considering in closer detail the motivations offered by the Relativist when it comes to defend Relativism on the account that it captures disagreement on such matters as taste, ethics and aesthetics. First of all, what is a disagreement? There is a very immediate notion of disagreement that the theorist has available when trying to figure out an answer to this question – whether or not in the area of personal taste. MacFarlane (2007b) calls it the "Simple View", according to which:

(D1) A disagreement obtains just in case there is a proposition *p* such that A accepts that *p* and B rejects that *p*.

or equivalently:

(D2) Disagreement obtains just in case A accepts *p* and B accepts *q*, whose truth implies *not-p*.

Formulations (D1) and (D2) both imply that A's acceptance and B's rejection of *p* (or A's acceptance of *p* and B's acceptance of *q*) are such that neither of the two parties could accept what the other accepts without becoming incoherent, or changing her mind on the matter: this is what MacFarlane (MS) calls "doxastic non-cotenability".

The Simple View on disagreement is exclusively focussed on disagreement as depending on an incompatibility between the contents accepted by the parties in the dispute, or at most on an incompatibility between attitudes (acceptance and rejection) which take as their object a common content. If one has already subscribed to the basic framework that underlies Relativism, however, it takes little reflection to realise that the incompatibility between contents, as well as the incompatibility between attitudes which take as object the same content, is neither necessary nor sufficient for disagreement (see MacFarlane 2007b, MS).

Suppose there are tensed propositions, i.e. propositions that are neutral with respect to times. Then the sameness of proposition accepted and rejected doesn't suffice for disagreement. Suppose that Mary accepts *<John is eating a sandwich>* at 2 pm and Claire rejects *<John is eating a sandwich>* at 3 pm. The tensed propositions that Mary and Clare respectively accept and reject are the same, though the situation doesn't qualify as a disagreement because the acceptance and rejection concern two different times (2 pm and 3 pm respectively). Of course, the point holds the same even if we represent the disagreement as Mary's acceptance of *<John is eating a sandwich>* at 2 pm coupled with Claire's acceptance of the contradictory sentence *<John is not eating a sandwich>* at 3 pm. The two acceptances take contradictory tensed propositions as objects, but disagreement doesn't obtain because the acceptances concern different circumstances (2 pm and 3 pm respectively).

Moreover, sameness of proposition accepted and rejected is not even necessary: suppose Mary

accepts *<John is eating a sandwich>* at 2 pm and Claire accepts *<Nobody was eating a sandwich one hour ago>* at 3 pm. On the assumption that the propositions are time-neutral, it turns out that the contents of the acceptance and rejection are different and even compatible (one could entertain both of them at the same time), though the situation constitutes a disagreement, because the acceptance and rejection concern the same time—viz. 2 pm.

Note that these considerations hold independently of whether one believes in tensed propositions. To demonstrate this, MacFarlane imagines the following case, in which the same classic proposition is accepted and rejected in two different worlds, and yet the situation doesn't amount to a disagreement (but see also Richard 2011 for a very similar example and considerations):

> Consider Jane (who inhabits this world, the actual world) and June, her counterpart in another possible world. Jane asserts that Mars has two moons, and June denies this very proposition. Do they disagree? Not in any real way. Jane's assertion concerns our world, while June's concerns hers. If June lives in a world where Mars has three moons, her denial may be just as correct as Jane's assertion. (MacFarlane 2007b: 23).

These remarks lean themselves to the view that whether two parties are in disagreement is not always a matter of whether they are accepting and rejecting the same proposition, but rather of whether the *circumstances* in which they are holding an attitude each with respect to some content are such that the correctness of the attitude of the first party at the first circumstance turns out as incompatible with the correctness of the attitude of the second party at the second circumstance, and *vice-versa*.

In order to characterise disagreement more precisely, MacFarlane (2007b, MS) introduces the notion of *accuracy*: an assertion is accurate iff its content is true at the relevant circumstance of evaluation (which in many cases is the circumstance of the context of use). Disagreement then can be pinpointed as an incompatibility (or "preclusion") of accuracy between two acceptances performed concerning certain circumstances of evaluation. MacFarlane terms the kind of disagreement exemplified as "preclusion of joint accuracy". The ensuing definition is as follows:

(D3) A disagreement obtains whenever A accepts that *p* and B rejects that *p* and A's acceptance and B's rejection cannot be both accurate.

MacFarlane then moves from disagreement about objective matters to disagreement about *subjective* matters. In the subjective area, which includes taste, ethics and aesthetics, his characterisation of disagreement works only if it is combined with the additional Relativist assumption that the preclusion of joint accuracy obtains relative to a *context of assessment*. So, on the assumption that Mary's assertion of "Ice cream is tasty" is true iff ice cream is tasty in the world w_u of the utterance and at the standard s_a of the assessor and that, analogously, Claire's assertion of "Ice cream is not tasty" is true iff ice cream is not tasty in the world w_u of the utterance and at the standard s_a of the assessor, then Mary and Claire are in disagreement insofar as Mary's and Claire's assertions cannot both be accurate with respect to the same (standard of a) context of assessment s_a. This way of describing the disagreement is taken to allow the Relativist to capture disagreement on disputes of inclination—as already explained also in **section 3**[35].

To summarise, MacFarlane's strategy consists in, firstly, defending a conception of disagreement as the one stated in (D3) and, secondly, in accounting for disagreement on matters of inclination by

---

[35] (D3) can be ameliorated—as MacFarlane himself recognises—so as to encompass cases in which A and B accept *different* contents, and yet the accuracies of their acceptances preclude each other. An improvement on (D3) is worked out by Rieppel (2011), who provides the following characterisation of disagreement: "A and B are in disagreement iff the fulfilment of the accuracy condition of A's acceptance guarantees the fulfilment of the inaccuracy condition of B's rejection, or vice versa." (Rieppel 2011: 251). As one can note, there's no mention in Rieppel's formulation of the fact that A and B are entertaining a common content *p*, contrary to what obtains in (D3).

combining (D3) with his notion of accuracy *as relative to contexts of assessment*. This second move in the strategy will prove important for our purposes: As we will see, MacFarlane's postulation of contexts of assessment is a legitimate move only if there's already the presumption that the semantics of the expressions at interest is sensitive to *"perspectival" matters* (e.g. tastes, moral codes, etc.). If this presumption is not in place—as I will argue to be the case with conceptual constraints—the strategy seems ultimately unmotivated.

The second leading motivation on which I would like to dwell corresponds to the Relativist's point to the effect that Relativism helps make sense of *retraction* in some key areas of discourse. As mentioned in **section 3**, retracting means, informally, rejecting an assertion that one has made previously (on the basis of some relevant reasons). The phenomenon of retraction is interesting for the Relativist as soon as one enters areas of discourse that are typically addressed as subjective. Here the phenomenon of retraction can help the Relativist to sharply trace the difference between context of utterance and context of assessment. Let us see how.

In general, a subject S asserts what S believes is true in the context occupied by S while performing the utterance. This means that, most of the times, the context c_a in which the subject assesses her own belief is the same as the context c_u in which she performs the utterance[36]. This fact about how we conduct our practices of assertion is unfavourable for the Relativist, because it doesn't make the difference between c_u and c_a fully apparent. However, MacFarlane (MS) believes that there is one practice whose normative structure clearly displays the difference between c_u and c_a: retraction. According to MacFarlane, the rule for retraction is the following:

[Retraction]
A subject S in context c_2 ought to retract an assertion (not yet retracted) made in context c_1 if that assertion is not true as performed in c_1 and assessed from c_2.

The norm mentions two contexts: c_1 is the context in which the assertion is performed; c_2 is the context in which the assertion is assessed. If there are examples of retraction that comply with this rule, i.e. if in linguistic practices speakers behave in a manner which can be regarded as responding to this rule, then the Relativist has made a strong case in favour of his view. *Prima facie*, it seems that retractions about taste can be made to fit within this template. For example, if at t_1 I utter "Fish sticks are tasty" on the basis of my taste at t_1, and then at t_2 I change my taste, I can felicitously retract my previous assertion by claiming: "I was wrong. Fish sticks are not tasty". In doing this, it seems that I am assessing my previous utterance from a different point of view—which in MacFarlane's framework is nothing but a "context of assessment". If this account of retraction on matters of taste is correct, then the Relativist has made a case for his theory. Once again, note that the phenomenon of retraction can provide support for the relativist only if retraction is cashed out in terms of contexts of assessment. This is possible in an area of discourse in which the presumption is already very strong that the truth of assertions in that area will be relative to "perspectival" aspects —e.g. tastes, aesthetic and moral standards. If this presumption falls apart, however—and I will argue that this is the case with conceptual constraints—an argument that draws on retraction doesn't really provide a motivation for embracing Relativism.

After having presented how the Relativist motivates her view by appealing to the way it captures disagreement and retraction in "subjective" areas of discourse, let us turn to the question *whether anything like Relativism is ever needed in order to account for the role of conceptual constraints* in determining the truth-conditions of an utterance of "It's raining", but also "Jill is ready", "Lisa is tall", "The leaves are green" etc. As mentioned in **section 2**, the basic move that would underlie any

---

[36] This is not always the case. There may be cases of "allocentric" assertion, in which the speaker asserts what she believes is true (or would be true) on another context of assessment. (see Lasersohn 2005, Stephenson 2007). But this happens in exceptional cases—e.g. when the putative assessor cannot speak for herself, as it happens with young children, or pets.

relativistic approach to the semantics of these sentences would imply countenancing an extra-parameter in the circumstances of evaluation for each of these sentences. As an operative proposal, I envisaged the truth-conditions of the previously mentioned sentences as being relativised to an extra-parameter besides the possible-world parameter, in the following way:

(10) "It's raining" is true $<w, l>$ iff it's raining at w, l.
(11) "Jill is ready" is true $<w, p>$ iff Jill is ready at w, p.
(12) "Lisa is tall" is true $<w, s>$ iff Lisa is tall at w, s.
(13) "The leaves are green" is true $<w, r>$ iff the leaves are green at w, r.

where l is a location parameter, p is a purpose-parameter, s could be regarded as a standard-of-height parameter, while r could be viewed as a parameter ranging over "respects". Note that I could have avoided the postulation of multiple kinds of parameters ranging each over a different feature by adding a single, catch-all "counts-as" parameter, like in MacFarlane (2007a, 2009). Since the outline of the details is not crucially important for our purposes, for the present moment let us stick to this variety of coordinates in the circumstances of evaluation. The question to be faced is the following: is it sensible to adopt a full-blown Relativist strategy, and therefore posit that the truth of utterances of "It's raining", "Jill is ready", etc. is true relative to a *context of assessment*?

The point I wish to make is that there is little use, from a purely theoretical point of view, for the notion of context of assessment on this occasion. In order to see this, it is convenient to make a conceptual distinction between what I shall call "perspectivality" and "feature-sensitivity". Perspectivality may be defined as relativity of truth to *ways of looking at the world* that cannot be reduced to aspects of reality itself: frames of reference, moral, aesthetic or taste stances, evidential bases, standards of precision and so on. "Feature-sensitivity" may be defined as relativity of truth to *aspects of the world*: I submit that these aspects include at least times and locations, but probably also comparison classes, contrast classes, schedules, parts, practical interests and stakes[37] [38]. This distinction is worth mentioning in that it helps to make clear that, arguably, there is an important difference between the truth-relativity of "It's raining" and its semantically under-determined cognates—in which we are presently interested in—and the truth-relativity of the sentences the Relativist is concerned with, like "Ice cream is tasty": the former is feature-sensitive, while the latter is perspectival. The truth of "Ice cream is tasty" is perspectival in that it is relative to a way of looking at the world—namely a gustatory stance, or standard. The truth of "It's raining" is feature-sensitive in that it is relative to an aspect of the world—namely a location. Analogously, the truth of "Jill is ready" is relative only in the sense that it depends on a feature of the world, namely a purpose or activity; while the truth of "The leaves are green" is relative in the sense that it depends on a respect of being green—which is an aspect of the world and *not* a way of looking at the world[39]. Interestingly, there is a sense in which even the truth of "Ice cream is tasty" may be

---

[37] I consider practical interests and stakes as "aspects of the world" because, even though they concern subjects, they are nevertheless ways of "cutting up reality". Practical interests and stakes only produce "focal effects" on features of the world—they do not constitute points of view irreducible to any aspect to the world.

[38] Note that the difference between perspectivality and feature-sensitivity cuts across the difference between an extra-parameters setup and a more traditional, e.g. hidden indexical setup. The feature-sensitivity of an expression in itself doesn't suggest that an approach in terms of hidden indexicality would be the best option to capture the semantics of such an expression. One could envisage an extra-coordinate in the circumstances of evaluation for feature-sensitive terms, as temporalist do with times and e.g. Lewis (1980) does with times and locations. Vice-versa, the perspectivality of the truth of a certain sentence in itself doesn't imply that an extra-parameter approach would be the best semantic option; additional considerations are needed to build up a case for this idea, mainly regarding the way speakers appear to use the expressions at issue, the semantic predictions they give rise to, etc. So the distinction between perspectivality and feature-sensitivity does not map onto a distinction between an "extra-parameters" approaches as opposed to a more traditional, single-parameter approach based on hidden indexicality.

[39] The case of "green" may be confusing, but we shouldn't be led astray by the fact of using an expression like "*respect* of being green". Being such only on the outer surface is definitely a "respect" in which an object can be

regarded as feature-sensitive: for example, it may be viewed as relative to individuals at times, or to standards set by some objective aspect of the world (experts, biological facts, etc.).

I contend that arguments from disagreement and retraction may establish Relativism in the way canvassed by MacFarlane only when so-called contexts of assessment are determined by *perspectival* aspects; when contexts of assessment are determined by features of the world (locations, times and so on) the Relativist's arguments are powerless. Since conceptual constraints determine aspects that could make truth at most *feature-sensitive*, adopting Relativism in order to account for conceptual constraints would be unmotivated. Let us see why this is so in greater detail.

In MacFarlane's account, the Relativist set-up in terms of contexts of assessment is particularly well suited to account for—among others—disagreements in the personal taste area. So if A asserts "Ice cream is tasty" and B asserts "Ice cream is not tasty", by the Relativist's lights the two parties are having a disagreement insofar as the accuracy of the assertion of the one is incompatible with the accuracy of assertion of the other in the same context of assessment. However, one could object, every factual disagreement is *also* such that there are two claims involved, which are incompatible with respect to one and the same context of assessment. For example, if A claims "Two plus two equals four" and B claims "Two plus two doesn't equal four", they certainly are in a disagreement insofar as the accuracy of A's claim is incompatible with the accuracy of B's claim with respect to the same context of assessment. However, this *by itself* doesn't open the way to Relativism. In order to establish Relativism, what is needed is the idea that there are many contexts of assessment, that all of them are legitimate, or on a par; and this idea is admissible only in a few cases, as e.g. in the case of taste, ethics, epistemic modality, future contingency..., that is in all those cases in which there's a previous, strong inclination in favour of a *perspectival* view of truth in the relevant area of discourse.

The intuition is, however, that the kind of truth-relativity that conceptual constraints give rise to is not perspectival, but rather *feature-sensitive*, i.e. it is relativity of truth to aspects of the world (locations, times, comparison classes, parts, etc.). This turns out as *incompatible* with the idea that we've seen as underlying the case in favour of Relativism, i.e. that the multiple contexts of assessment are all on a par. Let us see how this might become apparent. Suppose A, speaking about Paris, asserts "It's raining" while B, also talking about Paris, asserts "It's not raining". There is obviously an incompatibility between the accuracy of A's and B's respective assertions within the same context of assessment: not both of them can be accurate as statements concerning the weather in Paris. The problem is that there's no room here for advancing a Relativist account for this kind of disagreement, because clearly, even though A and B occupy two different "contexts of assessment", not both of them are equally legitimate. Suppose it's raining in Paris, so that A's assertion is accurate. This entails that B's assertion is inaccurate from A's context of assessment, but not only that: B's assertion is inaccurate from *all* possible contexts of assessment. That is because, *pace* the postulation of contexts of assessments, the accuracy of A's assertion is absolute and not relative. So, no favourable conclusion for Relativism is to be expected from cases of disagreement involving feature-sensitivity of truth. If the application of conceptual constraints gives rise to feature-sensitivity then, Relativism is not a good account for conceptual constraints. Adopting Relativism with respect to the semantics of sentences such as "It's raining", "Jill is ready", "The leaves are green" and so on would be ultimately unmotivated.

Similar concerns apply with respect to the Relativist's case from retraction. Let us make an example. Suppose that, on the basis of my evidence at t_1, I (believe and) assert "It's raining" while speaking about Paris. Shortly afterwards, at t_2, I discover that what I thought was rain was really water falling from some above apartment. As a consequence, I take my assertion back, since I

---

said green, and it is doubtless an aspect of the world. Being such under normal light conditions is another "respect" in which an object can be described as green, and it is once again an aspect of the world. No doubt a sentence like "This is green" could be true even with respect to a perspectival aspect, such as a subject's phenomenal experience. However, my concern is here that of emphasising the fact that a "respect" is not to be equated with a "perspective".

discover that, though it was supported by some evidence (and hence in a sense legitimate), my assertion at t_1 was nevertheless false. What is important here is that my assertion has always been false, even though I believed (perhaps justifiedly) that it was true. The world was arranged in such a way that the sentence "It's raining" as used to speak about Paris at t_1 was simply falsified by facts. The alleged "context of assessment" in which I perform the retraction is simply the context in which I discover that my assertion has always been absolutely false. In this case, however, the fact that I occupy a context of assessment from which I reject a previous assertion doesn't speak in favour of Relativism. It would speak in favour of Relativism if my assertion at t_1 was true with respect to t_1, while false with respect to t_2. But this is simply not the case. Since the locational aspects of the truth-conditions of my utterance of "It's raining" is a feature of the world and not a perspective, an argument from retraction which employs as example a retraction of an utterance of "It's raining" does not bring any support to Relativism. Thus, if cashing out conceptual constraints as parameters of evaluation for sentences like "It's raining" gives rise to feature-sensitivity, then Relativism shouldn't be adopted as an account of the semantics of these expressions.

The same reasoning applies to "Jill is ready", "Lisa is tall", "The leaves are green" etc. Let me briefly discuss the case of a sentence like "The leaves are green", about which one may object that its truth is not after all feature-sensitive—for it depends on "respects" of being green or, equivalently, on what "counts as" green according to the speaker's interests, where these notions strongly suggest a *perspectival* aspect. Suppose that Pia utters "The leaves are green" in context c_1, where being green on the outer surface is enough for something to count as green. It seems that "being green on the outer surface" is an aspect of the world (an objective property), and not a perspective. Suppose that Lia disagrees with Pia, and hence asserts "The leaves are not green", because she sees their surface differently. Now what settles the question between Pia and Lia is an objective aspect of the world—whether the leaves are or not green on the outer surface. There is no room for a relativistic account of disagreement in this case, because "respects" of being green (or ways of "counting as" green) do not determine perspectival aspects, rather they determine genuinely objective aspects. (Note in passing that sentences like "The leaves are green" *can* be evaluated also relative to perspectival aspects—such as perceptual experiences. So suppose Pia uttered "The leaves are green" in c_2, where "looking green on the outer surface" is enough to count as green. This is definitely a perspectival aspect. Could one construe an argument from disagreement to support Relativism on "green"? I doubt it, since if Lia uttered "The leaves are not green" on the basis of her perceptual experience, she would be right in her own perspective, and there would be no disagreement. A relativist argument from disagreement about "green" could not even get off the ground.)

To summarise, in this section I have set forth an argument which obeys to the following general schema:

1. Relativism is (allegedly) established by arguments from disagreement and retraction;
2. These arguments make Relativism motivated only if there is already a strong inclination to presume that truth in these areas of discourse is *perspectival*;
3. Conceptual constraints give rise to a view of truth as at best *feature-sensitive*;
4. As a consequence of this, there is no strong inclination to believe that the area of discourse in which the question of conceptual constraints emerges is one where truth is perspectival;
5. So Relativism would be unmotivated as an account of conceptual constraints.

After having ruled out Relativism as an account of conceptual constraints, what alternatives do we have left on the table? What distinguishes Relativism from Non-Indexical Contextualism is the postulation of contexts of assessment; given that the device of contexts of assessment was found defective for the purpose of laying down an account of conceptual constraints, a plausible alternative for the theorist would be to just drop contexts of assessment while still sticking to a

framework which countenances truth as relative to extra-parameters. This could bring one to embrace Non-Indexical Contextualism, which postulates extra-parameters but merely states that their value should be fixed at the context of *utterance*. In the next section, I shall deal with whether Non-Indexical Contextualism is enough motivated as well as an account for conceptual constraints.

5. Conceptual Constraints: Is Non-Indexical Contextualism an Option?

None of the formal semantics representations of conceptual constraints that we have covered so far, namely Indexical Contextualism (in **Chapter 3**), and Relativism (in **section 4** of the present chapter), seem to be good candidates for representing conceptual constraints.

As I have argued, the choice of Relativism isn't motivated enough in order to account for the conceptual constraints that seem to govern comprehension and evaluation of "It's raining", "Jill is ready", "The leaves are green" and so on. The reason is that Relativism seems to be appropriate as a semantic account of all those expressions whose truth is "perspectival", i.e. whose truth depends (or at least there's a strong intuition that it depends) on ways of looking at the world which cannot be reduced to any objective aspect of reality; these "ways of looking at the world" are multiple and all legitimate, and can give rise to equally accurate evaluations. Sure enough, the evaluation of an assertion of "It's raining" or "The leaves are green" might *as well* depend on perspectival aspects, as e.g. the standard of precision the speaker is currently adopting. But note that this perspectival aspect is not what is at issue in the characterisation of how conceptual constraints may make truth relative. In order to characterise the way conceptual constraints may determine truth-relativity, I spoke of "feature-sensitivity", i.e. relativity of truth to aspects of the world, as e.g. locations, times, comparison classes, parts. These "aspects of the world" are objective and subject-independent, and they determine the truth-value of an utterance in an eternal, stable way. It may be that predicates like "rain", "being ready", etc. are sensitive to both perspectival elements and to features of the world; this however, only implies that the same expressions show more than one type of semantic context-sensitivity (for a version of the idea of a double context-dependency in semantics, see Lasersohn 2011). What interests us is that peculiar kind of context sensitivity which makes the semantics of these expressions sensitive to elements whose specification determines *how the world should be* in order for the expression to have a truth-value, as opposed to how the world should be *from some point of view*.

Given the putative greater appropriateness of the notion of feature-sensitivity to describe how conceptual constraints might operate on the (technical) notion of truth in order to overcome the specific problem we are trying to solve, namely that of semantic under-determinacy, it seems worth exploring the hypothesis that the best option for the theorist who wishes to cash out conceptual constraints as evaluation parameters would be to adopt *Non-Indexical Contextualism* as opposed to Relativism as Assessment Sensitivity.

The question that becomes pressing at this stage is therefore the following: are there good reasons to adopt Non-Indexical Contextualism as an account of conceptual constraints?

To begin with, considerations of *elegance and parsimony* may be invoked. Non-Indexical Contextualism, it may be said, regards expressions such as "to rain", "being ready", "being green" as being invariant, non-relational and non-indexical, thus allowing the theorist to explain intuitions in a host of cases. Here are some examples.

- *Speech Reports:* Holding that "ready" is a non-relational, non-indexical expression allows us to see why, if A asserts "X is ready" in context c_1, B can easily and accurately report A's assertion by uttering "A said that X is ready" in context c_2. The reason is precisely that the report is accurate because the reported predicate "ready" in c_2 is the same as the predicate "ready" used in c_1;
- *Same-sayings:* Non-Indexical Contextualism has a straightforward explanation for why one has the intuition that, if A utters "X is ready" in c_1 and B utters "X is ready" in c_2, A and B

"said the same thing". A and B literally said the same thing in that they used the very same monadic, invariant predicate;

(As an aside, note that curiously enough (and as we mentioned in **section 3**) these are the same advantages advertised by proponents of Invariantism, such as the Minimalists Cappelen and Lepore (2005) and Borg (2004), to support their own position. Just like Non-Indexical Contextualists, Invariantists claim that predicates like "ready" are monadic, so that intuitions on speech reports and same-sayings are accounted for in the Invariantist framework as well should be no surprise. The only difference between Invariantism and Non-Indexical Contextualism is that Invariantism does not espouse any form of extra-parameters relativity of truth, but sticks to an *absolute* notion of truth. This however, makes the position highly contentious, for it then becomes unclear how a sentence like "Jill is ready" or "It's raining" can express a truth-condition, once "being ready" or "to rain" are not only conceived as a monadic predicates, but also the truth value of the sentences is conceived as absolute. The critic can legitimately wonder in which case it is true that Jill is ready *tout court* in an absolute way, or in which case it is true that it rains *tout court* in an absolute way[40].)

The examples listed above all largely succeed in virtue of their use of a notion of "said that" which keeps track only of the "strict" semantic features of the expressions in question. However, intuitions might be reversed as soon as one sets up the scenario in the appropriate way, so as to put emphasis on the utterance's overall significance in context rather than on the sentence's strict semantic properties. Since the cases presented above are both about uses of a sentence that take place in two distinct contexts c_1 and c_2, I submit that providing details about these different contexts would make it more likely that one judges that the speakers across different contexts *do not* say the same thing. Consider the following scenario (the example is taken from Leslie 2007):

> C_1: John is woefully unprepared for his APA interviews, and is about to pass out from nerves. He really needs more time before his interviews. He states "I'm just not ready".
> C_2: There is an evangelist hovering around the convention, looking to see whether there are any lost souls for whom the time is right to accept Jesus Christ as their personal saviour. Spotting John, the evangelist says "Ah, he's ready!".

Given the larger amount of details provided, I take it that one's intuitions become at least less clear when it comes to apply the tests illustrated above about speech reports and same-sayings. First of all, one might not be entirely sure now whether John, *given the way he used the predicate "ready" in c_1*, could felicitously disquotationally report the evangelist by uttering "The evangelist said that I am ready", nor state anything like "When the evangelist says "He's ready", what he says that I am ready". Neither it is so sure that, in context c_3, a third party could felicitously state, in light of the

---

[40] The main reply to this objection, which can be found in Cappelen and Lepore (2005: 157-175) but also in Borg (2009: 48-9; 2012: 524) goes as follows: the critic's point is simply that minimal propositions such as *<that it's raining>* do not describe any state of affairs such that it could make them true. However, Minimalists note, this is a problem which obtains independently of whether one believes in minimal propositions. Even the proposition *<that it's raining in Paris>* may give no sufficient conditions for it to be true. After all, one may wonder: is it raining in Paris if it's raining only on the 6th arrondissement, or if it's raining only on the Northern part of the city—and so on? No sentence (and the related semantically expressed proposition) fully specifies what has to be the case in the world for it to be true. From this it is derived that, if even semantically *determined* sentences fail to specify how the world should be for them to be true, then the problem doesn't threaten minimal propositions. After all, even minimal propositions like *<that it's raining>* are such that they determine a set of possible worlds where these propositions are true—namely all those worlds in which it's raining. However, one might observe, all worlds where it's raining are worlds where it's raining *somewhere*. There is no set of worlds where it's raining *tout court*. So there can't be minimal propositions *<that it's raining>* sorting out worlds where it's raining, full stop. Either we admit for states of affairs of it being raining *tout court* (which is absurd) or we admit that so-called "minimal propositions" are more fine-grained (and less minimal) than minimalists want.

different contextual details: "John and the evangelist said the same".

(Note that, in contrast with the Invariantist/Minimalist, who can explain these contrasting intuitions by appealing to the fact that the *speech acts* performed can be taken to express different (non-minimal) propositions (see for instance Cappelen 2008), the Non-Indexical Contextualist doesn't have the means to explain intuitions that involve attributions of different *contents* to the two assertions for, in his view, the contents of the assertion never differ—what differs is rather the circumstances of evaluation and the truth-value. So intuitions as to the variability of content remain unexplained in this framework.)

The upshot of these considerations is that building up a case that strongly relies on intuitions about "what is said" is not an optimal strategy. Intuitions on who said what are unstable and subject to the presence of background conditions and details: I take it that the case for Non-Indexical Contextualism is considerably weakened by the possibility of going back and forth between contrasting verdicts as to what the speakers "said".

The same considerations could be advanced also with respect to tests that employ reports in terms of "agree/disagree". The Non-Indexical Contextualist could claim that the following is an explanatory advantage of her view: If A asserts "Jill is ready" in context c_1 and B asserts "Jill is ready" in c_2, C could consistently say "A and B agree that Jill is ready", because A and B are using the same monadic predicate "ready". But do they really *agree*? It depends on what A and B are talking about in their respective contexts (or, put in Austin's (1950) terms, it depends on which situations they are demonstratively referring to). The case becomes even less convincing in the case of disagreement: A asserts "Jill is ready" in context c_1 and B asserts "Jill is not ready" in c_2. Could C consistently say "A and B disagree over whether Jill is ready"? Well again, it depends on what A and B are talking about. Note that the case doesn't gain in plausibility even when we introduce disagreement on matters of personal taste. A asserts "Ice cream is tasty" in context c_1 and B asserts "Ice cream is not tasty" in c_2. C says "A and B disagree over whether ice cream is tasty". But do they? This is dubious even by the lights of Non-Indexical Contextualism: if each speaker's assertion is regarded as correct with respect to the speaker's own taste-standard, the sense of disagreement vanishes (see Stojanovic 2007 and MacFarlane 2007a for considerations in this vein). General considerations on taste disputes such as those offered by Cappelen and Hawthorne 2009 may add up. Cappelen and Hawthorne propose a "sceptic therapy" when confronted with an apparent taste-disagreement: if, in describing the situation, we take factors such as physical dispositions, character differences, culture and upbringing into consideration, we may be persuaded that no talk of disagreement is properly in place here.

The advantages offered by Non-Indexical Contextualism so far don't look as exactly sharp, in that the intuitions that Non-Indexical Contexualism wishes to exploit are not sharp either, to the extent that they employ the rather unstable notion of "what is said", or "agree/disagree". But the troubles for Non-Indexical Contextualism are not over. It has been recently pointed out (see Kissine 2010) that Non-Indexical Contextualism gives wrong predictions of compatibility between the contents of assertions in scenarios like the following:

(14) Alvin, speaking about John's preparation for the APA interviews: "John isn't ready"
      The evangelist: "I agree. John isn't ready for the APA interviews, though he's *ready*".

Let us assume, for ease of exposition, that Alvin's assertion is true at a circumstance of evaluation consisting of a possible world w and Alvin's purpose-parameter p. Now if the evangelist's expression of agreement is correct, this means that Alvin and the evangelist are evaluating their utterances from the same "point of evaluation" (represented as $<w, p>$). This, however, makes the evangelist's assertion of "He's ready" false. So the evangelist could not agree with Alvin in the sense of co-habiting the same point of evaluation and at the same time assert that John is ready. The puzzle seems to arise from the fact that the Non-Indexical Contextualist (just as the Relativist in this

respect) is committed to a context-sensitivity that is *global*, i.e. it affects an entire sentence, and not *local*, such as to affect parts of a sentence, like terms or phrases. Given the "global" nature of context-sensitivity characterising a Non-Indexical Contextualist framework, one cannot just take a conjunction "P & Q" and evaluate the conjuncts separately, at different circumstances of evaluation. This global feature of context-sensitivity gives rise to the puzzle: the evangelist cannot state that John isn't ready (for the APA interviews) and is ready (to be saved), for the context-sensitivity of his assertion ties the truth value of the sentence he uses to a single circumstance of evaluation, where the conjuncts cannot both be true. Note that Indexical Contextualism, which adopts a notion of context-sensitivity as *local*, i.e. which countenances the possibility to contextually affect the content of single sentence components (predicates, phrases), can indeed depict the evangelist as asserting *that John isn't ready for the APA interviews and is ready to be saved*, and therefore explains why the evangelist's claim "I agree" is appropriate in the context of the conversation. This further point seems to make Non-Indexical Contextualism less desirable as an account of conceptual constraints, for it fails to account for certain very strong intuitions about the way language is appropriately used. A last point that could be raised against Non-Indexical Contextualism has to do with the fact that, though one may regard it as a way of dealing with the problem of semantic under-determinacy, once the account is adopted semantic under-determinacy itself gets erased, so to speak. This is because Non-Indexical Contextualism regards sentences such as "It's raining", "Jill is ready", "The leaves are green" and so on as semantically *determined* as far as their content is concerned. The intuition of semantic under-determinacy though can be explained away by the Non-Indexical Contextualist framework insofar as, though it accepts that the sentences at interest are semantically determined as to their content, they are underdetermined as far as the circumstances in which they should be evaluated.

This way of looking at semantically under-determined expressions, which essentially involves explaining the intuition of content under-determinacy away, ultimately betrays the strong and—seemingly—well-rooted intuition that both theorists and ordinary speakers have, to the effect that these sentences *need completion at the level of what is said*. If Non-Indexical Contextualism were true, it would remain unexplained why the normal stance with respect to a sentence like "It's raining" is to claim that it "doesn't say enough". The Non-Indexical Contextualist could always invoke the hypothesis that language users are, so to speak, blind to the specific kind of semantic under-determinacy that the sentence really displays. This line of response however places a further burden onto the Non-Indexical Contextualist's shoulders: for now she owes an explanation of where this kind of blindness comes from. Do ordinary speakers—and a great deal of philosophers too—have some deeply entrenched bias? Are the intuitions of these subjects not enough "educated", perhaps naïve and in need for refinement? If this is so, however, is it fair for a theory like Non-Indexical Contextualism to dictate what our intuitions should be about?

The last question seems to me to raise a methodological problem, which has to do with the priority of data with respect to the theory. One would think that theories which take initial data seriously should not end up saying that, in light of the theory itself, the initial data were misguided. This is tantamount to denying data their *status* as data, departing from them in such a way that the theory goes as far as "trumping" its own initial data. I believe that the fact that a theory may lead to such a radical reconsideration of its initial data casts a dubious light on the theory itself (see Coliva and Moruzzi (MSa) on this point). In the present case, the fact that Non-Indexical Contextualism may lead to a radical reconsideration of intuitions concerning semantic under-determinacy—by saying that they are deeply misguided in that they don't really concern content, but circumstances of evaluation, seems to me to cast a non-favourable light onto Non-Indexical Contextualism itself.

To summarise the last argument: Evidence tells us that the language users' semantic intuitions as to the under-determinacy of "It's raining" *et alia* point in the direction of a defectiveness of *content*, and not in the direction of a defectiveness of circumstances of evaluation. Since Non-Indexical Contextualism implies precisely that intuitions as to content defectiveness are wrong—or "blind"

with respect to how things really are—this entails that Non-Indexical Contextualism, as a theory aimed at dealing with semantic under-determinacy, ultimately betrays its initial evidence. This (together with the considerations advanced in the previous paragraphs) constitutes a quite strong motivation for not adopting Non-Indexical Contextualism when it comes to capturing the phenomenon of semantic under-determinacy *as it is intuited by competent language users*.

6. Conceptually Mandated Content

So far I have been dealing with ways in which conceptual constraints could be implemented at the level of a formal semantic system. Conceptual constraints are, as I argued in **Chapter 2**, cognitive/discoursive configurations which may help hearers in the comprehension of semantically under-determined sentences such as "It's raining", "Jill is ready", "Lisa is tall", "The leaves are green" etc. In **Chapter 3**, I examined the possibility of regarding conceptual constraints as nothing but hidden indexical elements in the syntax of expressions, thus following what is known as Indexical Contextualism; in **section 4** of this chapter, I engaged with the question whether conceptual constraints could be cashed out in terms of Relativism, i.e. in terms of dependence of truth on so-called "contexts of assessment"; in **section 5** I focussed on the possibility that conceptual constraints be represented as simple extra-coordinates in the index of the circumstances of evaluation, whose value is fixed by the context of utterance (rather than by the context of assessment), in the vein of Non-Indexical Contextualism. None of the options I have considered proved to supply an adequate account of conceptual constraints at the level of formal semantics. So perhaps conceptual constraints should not be implemented at the formal level—be it the syntactic level, or the level of truth (metalinguistic level). My suggestion at this point is to look at conceptual constraints as operating *outside of the formal semantic machinery*.

At the basis of the account I will outline in this section is a view of the role that conceptual competences and constraints have in communication. We may start to ask what may connect a speaker's use of a semantically under-determined sentence with a hearer's understanding of it. I propose the following answer. The "bridge" between the utterer's use of a semantically under-determined sentence and the hearer's understanding of that utterance is precisely a(n implicit) presupposition of their *common conceptual (semantic) competences*. On the one hand, a competent speaker of English chooses to use under-articulated sentences ("It's raining", "Jill is ready") because: (a) she trusts that the hearer is as competent as the speaker about the English language and about the concepts invoked by the use of the English words that she uses; (b) she relies on the fact that, with the help of contextual clues and of conceptual/semantic constraints, the hearer will be able to grasp the content of the speaker's utterance. On the other hand, hearers who are in favourable conversational conditions, which means that they are recognised as participants to a conversation and know what the speaker is talking about, seldom question the use of under-determined sentences precisely because, with the help of the information they can gather from the conversational setting *plus* the application of the conceptual/semantic constraints they are familiar with, they are able to understand and evaluate what an utterance of an under-determined sentence says—on the tacit assumption that speakers are as competent as them about the language used and about the related conceptual constraints. The "bridging role" of conceptual constraints in verbal communication *via* semantically under-determined sentences could be thus spelled out:

> [The Bridge Thesis]
> In the production of semantically under-determined sentences, (a) speakers rely on the fact that their hearers (who are presumed as conceptually/semantically competent as them) will apply conceptual constraints (*plus* contextual information) in order to figure out the utterance's content and truth-conditions; and conversely (b) lingustically and conceptually competent hearers located in favourable conversational conditions will be ready to apply conceptual constraints in their

interpretation of these semantically under-determined utterances (trusting that speakers are as competent as them).

Having laid down the role of conceptual constraints in communication *via* utterances of semantically under-determined sentences, let me move on to a characterisation of the *status* of these constraints and of the result of their application.

To begin with, I take conceptual constraints to be *ways of organising our thought* so as to be able to effectively accomplish various cognitive tasks, as e.g. categorisation, memorisation, recognition. Conceptual constraints are primarily constraints on thought. However, conceptual constraints *de facto* coincide with *semantic* constraints on the use of linguistic expressions for all those speakers who are linguistically competent: this is because knowing the meaning of an expression *e* in a language *L* usually entails also being competent about the concept *F* which is associated to *e*. However, as I made clear in the **Introduction** and in **Chapter 2** (**section 10**), conceptual constraints may be independent of meaning constraints, for a subject could be competent about a concept *F* but be incompetent about the meaning of the expression *e* which, by hypothesis, is associated with *F* in *L*. The domain of conceptual constraints therefore may overlap with the domain of semantic constraints for those speakers who are competent users of a certain expression *e* of *L*, even though it is conceivable that one is conceptually competent on *F* without being competent on *e* (which is by hypothesis associated with *F* in *L*).

Conceptual/semantic constraints help enrich the content of certain sentences—in the present case, we are focussing on semantically under-determined ones. So the conceptual constraint that every event of raining is spatially located—corresponding to the semantic constraint that "It's raining" is correctly used iff it's raining at a location—may help the interpretation of utterances of "It's raining" in that it helps "expand" the under-determined semantic content *<that it's raining>* which results from compositional, syntax-driven interpretation only. The upshot is something like the following picture: for an utterance of "It's raining", compositional, syntax-driven semantics delivers the under-determined content *<that it's raining>*; conceptual/semantic competence enables the hearer to expand this defective content by adding a locational component, so as to obtain *<it's raining in x>*. Contextual information may be plugged in so as to assign a value to the variable *x* which occupies the conceptual/semantic slot for locations. If the salient location in context is Paris, the result might be *<that it's raining in Paris>*. If no salient location is available, the location variable may be bound by an existential quantifier, the result being *<that it's raining somewhere>*. As one may appreciate, conceptual/semantic constraints operate on the result of formal, compositional semantic processes to deliver an "expanded" semantic content. So on the one hand, we have formal, compositional semantics, which responds to syntactic/ lexical constraints and delivers a minimal, often semantically under-determined content; on the other hand we have constraints (at the conceptual/semantic level) that are "semantic" in a broader sense of the term: they pose restrictions on correct use of words and on "what is said" with an utterance of a sentence, but they do not do that as a result of any syntactic/lexical rules. The present picture therefore opens the way to *semantic processes which are syntactically unconstrained*. The following schema captures the relation between semantics, syntax and conceptual/semantic constraints.

Semantics (concerns correct use of words and the truth-conditions of utterances)	Syntactic/ lexical constraints
	Conceptual/semantic constraints

If there are semantic processes which nevertheless do not respond to syntactic constraints, then the propositional components which are introduced via conceptual/semantic constraints are syntactically not represented—they are *unarticulated constituents*. The last remark could suggest that, since the result of the application of conceptual constraints involves the presence of

unarticulated constituents, this kind of content will not differ from what has been called *explicature* (by Carston), or *impliciture* (by Bach). My contention is that the difference between clear cases of explicature/impliciture and conceptually mandated content is that, while explicatures/implicatures are worked out for context-dependent, purely pragmatic reasons, conceptually mandated content is worked out for context-independent, ultimately semantic reasons (in the broadened sense of "semantics" characterised above). Let me explain this distinction.

Consider how a hearer may be regarded as working out the content expressed by the speaker with an utterance of any of these sentences:

(15) Jack and Jill got married
(16) You are not going to die
(17) It will take time to get there
(18) I've had breakfast

Doing the formal semantics for these sentences supposedly delivers contents which can receive an evaluation in context. I take (15) to be true in a context iff Jill got married and Jack got married; (16) is true iff the addressee is not going to die in the future; (17) is true iff a certain trip, which is supposedly referred to with the pronoun "it", is going to take a certain amount of time to be completed; (18) is true iff the speaker has had breakfast in the past. In all these cases, the semantic content obtained compositionally via a formal semantic machinery is evaluable in a context, so no "deep" semantic under-determinacy affects these sentences (for the distinction between "deep" and "shallow" semantic under-determinacy, see **Chapter 2, section 5**). The content obtained via formal semantics, however, fails to determine what the speaker usually means with an utterance of (15)-(18). Thus, with (15) one could mean that Jack and Jill got married *to each other*; with (16) one could mean that the addressee is not going to die *from a minor cut of her finger*; with (17) one could mean that *it will take more time than expected* to get there; with (18) one could mean that one has had breakfast *that very morning*. In all these cases, I take it that what guides the hearer in the working out of the communicated contents of (15)-(18) is context-dependent, purely pragmatic considerations as to what would make these contents maximally relevant, informative, etc. Hearers will supposedly search for the content which makes the utterer's contribution maximally relevant, useful, informative for the sake of the conversation. Examples of "explicatures" and "implicitures" like (15)-(18) are then worked out for *purely (and genuinely) pragmatic reasons* (here I am momentarily merging the two notions, even though their proponents consider them relevantly different[41]).

There's a significant difference between the processes that are regarded as resulting in explicatures or implicitures like (15)-(18) and the processes that—if the present account is on the right track—yield a content that is "conceptually mandated". As I said, clear cases of explicature/impliciture like (15)-(18) are wholly pragmatic in nature, in the sense that these levels of content are worked out for genuinely pragmatic reasons. There is, however, a range of expressions for which it doesn't seem entirely correct to claim that it is *purely* pragmatic processes that guide their contextual interpretation. These are the cases of "deep" semantic under-determinacy I have been concerned with so far, such as "It's raining", "Jill is ready", "Lisa is tall", "The leaves are green", etc. In all

---

[41]   The explicatures theorised by Carston are obtained from an initial propositional schema through processes that are guided by a principle of relevance. That is, the sentence's content-in-context is the (schematic) input of an inferential process that aims at fleshing out that content, maximising its relevance for the purposes of the conversation by at the same time saving the most cognitive effort. According to Carston, explicatures still belong to the domain of semantics, in that they correspond to what the hearer "says" rather than "implicates". Implicitures—as Bach calls them—differ from explicatures in that they are obtained through clearly "Gricean" processes that, in Bach's view, operate at an implicit level. Thus, given an utterance of a semantically under-determined content, the hearer is supposed to avail herself of the principle of cooperation and of the maxims in order to figure out what the speaker is saying at the implicit level. In Bach's view, implicitures belong to the domain of pragmatics.

these cases, it seems that purely pragmatic considerations are only partially involved. For these expressions, what really is needed is that they be completed in such a way as to result in evaluable contents. Before pragmatic, context-dependent considerations step in, these expressions need to be made "semantically kosher", i.e. they need to be interpreted in such as way that they can receive an evaluation. For example, the semantic structure of the sentence "It's raining", which by the lights of formal semantics is simply *<that it's raining>* needs to be expanded so as to include a locational aspect, if any propositional item is to be obtained. Analogously, the semantic structure of "Jill is ready", is to be expanded so as to make room for a purpose-aspect with respect to which Jill can be said ready or not ready. What helps in "expanding" the semantic structure of these sentences is *not* context-related, purely pragmatic considerations; it is rather *context-independent, purely conceptual (semantic) considerations* as to what it is to be raining, to be ready and so on. So the interpretation of utterances of "It's raining", "Jill is ready" seems to need that conceptual/semantic considerations be run *prior* to any purely pragmatic considerations as to relevance, informativeness, perspicuity, etc. The difference between cases of explicature/impliciture like (15)-(18) and cases in which semantic considerations are prior to purely pragmatic ones may be cashed out as a difference between explicatures/implicitures and instances of *conceptually mandated content*.

Explicatures/Implicitures	Conceptually mandated content
"Jack and Jill got married" "You are not going to die" "It will take time to get there" "I've had breakfast"	"It's raining", "Jill is ready", "Lisa is tall", "The leaves are green".
Context-dependent, purely pragmatic considerations guide interpretation.	Context-independent considerations are prior with respect to purely pragmatic considerations.

One may point out that it's *just wrong* to think that utterance interpretation may be accomplished only with the help of conceptual/semantic competences, for conceptual/semantic competences may not help at all to find out what the speaker means in a context: In order to figure that out, we certainly need some pragmatic mechanisms! The objection is well taken. Of course, in the working out of an utterance's "conceptually mandated" content in the specific context in which the utterance is performed, the hear may need to retrieve information which is specific to the context of utterance. So for example, in working out the content of A's utterance of "It's raining" while speaking of location X, B certainly needs to find out information about X, and this can be done only by looking at the context. Here pragmatic considerations may well play a role: I don't wish to deny that pragmatic principles may accompany the application of conceptual constraints, though I wish to maintain that pragmatic considerations will always aim at identifying *elements of the background* such as "what the speaker is talking about", "what the topic of the conversation might be", etc. Pragmatic considerations may play a role in establishing issues as to what the referential intentions of the speaker are, or as to what topic in the conversation she is currently addressing, and so on. But note that these pragmatic considerations may only be employed in order to "fill the gaps" that conceptual/semantic constraints open as a result of their application, and not to *create* those gaps. Purely conceptual considerations may elicit the location-slot at the level of conceptual/semantic structure. This slot having been created on the basis of context-independent information, pragmatic considerations as to "which location the speaker is talking about" may provide location X (e.g. Paris) as the element which is most suitable to "fill the conceptual gap" just created by the application of one's conceptual competence. The following picture represents the structural difference between processes driven by conceptual constraints as opposed to processes driven by purely pragmatic considerations.

Conceptual Constraints View	Pragmatic Enrichment View
Input: Semantic content-in-context	Input: Semantic content-in-context
Application of conceptual constraints	Pragmatic considerations (Relevance, Cooperation)
Contextual info "fills in" conceptual gaps	
Output: Utterance content	Output: Utterance content

A very straightforward objection to the proposal may be the following: What is the advantage of adopting a conceptual constraints view as against a pragmatic enrichment view? The latter seems much simpler, while the former looks more complicated. Why should we suppose that processes of utterance comprehension of "It's raining", "Jill is ready" etc. are this complicated? As a first reaction, let me express my skepticism as to the conclusiveness of considerations concerning "simplicity" and "complication". That a process is represented as "complicated" doesn't necessarily entail that that process is more costly in terms of cognitive effort. In particular, when we bring conceptual constraints into the picture, chances are that their application makes the processes at issue even faster and less effort-taking in virtue of the fact that conceptual constraints may be thought of as designed to pre-organise thought in a very compact way. Be that as it may, the objection fails to appreciate the fact that the conceptual constraints view is not altogether incompatible with a pragmatic enrichment view; indeed, I am inclined to think that the conceptual constraints view may be regarded as a way of *clarifying and refining* a pragmatic explanation of the comprehension of at least certain sentences—namely the "deeply" semantically under-determined sentences that we have been interested in so far. A conceptual constraints view seems to account very well for the fact that the pragmatic enrichment of utterances of "It's raining", "Jill is ready" etc. *is not semantically unconstrained*, as opposed to the pragmatic enrichment of e.g. non-literal discourse (metaphorical, metonimical talk), or of sentences like (15)-(18). So, all in all, the greater "complication" introduced by the conceptual constraints view may not be a theoretical loss, but a theoretical gain, in that it doesn't necessarily entail any additional cognitive efforts and it explains the discrepancy between purely pragmatics-guided utterance interpretation (such as that found in non-literal talk and in (15)-(18)) and the apparently more constrained interpretation of "It's raining", "Jill is ready", "The leaves are green" *et similia*.

To sum up, in this section I have proposed to regard conceptual/semantic constraints as operating *outside of the formal semantic machinery*. Conceptual/semantic constraints act as a "bridge" between the utterance of an under-determined sentence and its interpretation, in the sense that speakers who are conceptually/semantically competent choose to utter such sentences relying on the fact that hearers, who are presupposed as equally competent, will be able to reconstruct the content of their utterances by employing such constraints. On the other hand hearers accept utterances of semantically under-determined sentences because they trust the fact that speakers are conceptually/semantically competent and expect them to employ the same competences to work out their content. Semantic/conceptual constraints give rise to semantic processes of "content expansion" which are nevertheless syntactically unconstrained. So the conceptual constraints view allows for elements of the proposition semantically expressed which nevertheless do not correspond to any articulated items in the sentence's syntax. The presence of unarticulated constituents in the proposition worked out by the hearer should not lead one to conclude that contents obtained through the exercise of one's semantic/conceptual competences are equivalent to explicatures/implicitures like the ones in (15)-(18): the latter are determined by context-specific, purely pragmatic considerations as to relevance, informativeness, perspicuity, etc., while the former are determined

on the basis of context-independent, semantic/conceptual considerations as to what makes uses of a certain concept/expression correct and veridical (*plus* contextual information, where appropriate).

Conclusion

In this chapter, I have considered the issue whether conceptual constraints should be regarded as parameters to which truth is made relative. This strategy would: (i) avoid commitments as to the presence of covert elements articulated in the syntactic structure of the sentences under examination (unlike Indexical Contextualism); (ii) account for the fact that speakers know at least the intuitive truth-conditions of utterances of under-determined sentences when these are listened "out of context" (unlike Radical Contextualism).

- As a first step, I have introduced the "basic move", which consists in introducing an extra-parameter in the index of the circumstances of evaluation;
- I have illustrated how the truth-conditions of semantically under-determined sentences ("It's raining", etc.) may be modified after the introduction of extra-parameters in the circumstances of evaluation;
- I have introduced Non Indexical Contextualism and (Radical) Relativism, two semantic approaches that share the "basic move". Non-Indexical Contextualism has it that truth is relative to extra parameters (besides a possible world parameter) and that the value of such parameters is fixed at the *context of utterance*. (Radical) Relativism has it that truth is relative to extra parameters (besides a possible world parameter) and that the value of such parameters is fixed at the *context of assessment*;
- I have then considered the question whether Relativism is enough motivated as an account of conceptual constraints. The most interesting defences of Relativism argue that Relativism accounts particularly well for *disagreement* and *retraction* in areas of discourse such as taste, ethics, epistemic modality. I argue that Relativism may successfully account for disagreement and retraction only when the presumed "context of assessment" to which truth is made relative corresponds to a *perspective*, i.e. a way of looking at the world that doesn't correspond to any aspect of the world. However, conceptual constraints do not give rise to perspectives, but only to objective features of the world, as e.g. locations, times, comparison classes, parts. This implies that Relativism is not ultimately motivated as an account of conceptual constraints;
- I have then moved on to considering whether Non-Indexical Contextualism could capture the role of conceptual constraints better than Relativism. I have shown that Non-Indexical Contextualism may not be an optimal choice in order to account for conceptual constraints, because of problems that theorists have emphasised in recent work. Most and foremost, though, Non-Indexical Contextualism is not desirable in that it imposes a radical revision of the nature and content of intuitions of under-determinacy, a result which I have argued to be methodologically suspect;
- The conclusion of the foregoing considerations is that the semantic strategies so far scrutinised—Indexical Contextualism, Relativism and Non-Indexical Contextualism—are all inadequate in order to capture conceptual constraints. *My proposal is to regard conceptual constraints as operating outside the formal semantic machinery;*
- I propose to regard conceptual constraints as ways of organising thought (and discourse) that hearers may employ to figure out the content and truth-conditions of certain under-determined utterances. Speakers may choose to utter under-articulated sentences precisely because they rely on their hearer's capacity to exploit their conceptual (semantic) competence *plus* contextual information in order to figure out what they are saying;
- The application of conceptual/semantic constraints gives rise to semantic processes which are nevertheless *syntactically unconstrained*. So the conceptual constraints view allows for

elements of the proposition semantically expressed (in a context) which nevertheless do not correspond to any articulated items in the sentence's syntax. This notwithstanding, "conceptually mandated" content doesn't correspond to cases of explicature or impliciture like the ones in (15)-(18), in that the latter are the result of context-specific, purely pragmatic considerations, while the former are the result of context-independent, semantic considerations (which may be integrated with contextual information, where appropriate).

Chapter 5

Conceptual Constraints, Content and Context

In the last three chapters, I have addressed the question of how speakers manage to understand and evaluate utterances of "It's raining", "Jill is ready", "The leaves are green", even though the semantics of these sentences under-determines the content of their utterances. In **Chapter 2** I have argued that the comprehension of this group of sentences depends in good part on hearers exercising their *conceptual (semantic) competence*. The linguistic under-articulation typical of semantically under-determined sentences leaves out certain components of the proposition expressed by the utterance (the location of the rain, the purpose for being ready, the respect in which the leaves are green), however these components are sufficiently easy to retrieve for the audience to the extent that they may correspond to conceptual/semantic aspects that the hearer is already familiar with. The hearer may use such constraints in order to expand the conceptual/semantic structure of the sentence, and plug in contextual information (*Paris*, *skiing*, *on the outer surface*) where appropriate.

In **Chapter 3**, my aim has been that of showing that conceptual constraints are not amenable to be reduced to hidden syntactic elements (in the way of Indexical Contextualism), while in **Chapter 4** I offered reasons for not reducing conceptual constraints to coordinates for truth-relativisation— whether contexts of assessment (in the way of Relativism) or extra-parameters in the index of the circumstances of evaluation (in the way of Non-Indexical Contextualism). My suggestion, eventually, was to *regard conceptual constraints as operating outside the formal semantic machinery*. The idea, in a few words, is that the application of conceptual constraints is a process which is run in addition to strict, compositionally driven semantic interpretation. Applying conceptual constraints to sentential, compositionally obtained meaning implies supplementing it with conceptual/semantic aspects that are not encoded syntactically or lexically. Ultimately, the exercise of one's conceptual competences is complementary to the strict, compositionally driven semantic interpretation. This is, very sketchily, the proposal which I am going to develop in closer detail in the present chapter.

1. The Scope of the Proposal

In this section, my concern will be exclusively that of setting the boundaries of the proposal, in order to prevent as much as possible any misunderstanding as to the terrain on which I am operating. What I have been trying to explain throughout the last three chapters is the success of communication between language users when what is employed is *fully sentential*, though *non-proposition-expressing* linguistic entities. The examples I have been using all along are "It's raining", "Jill is ready", "Lisa is tall", "The leaves are green" etc.; all of them meet the requirement of being fully sentential constructions, which however do not, in a number of relevant cases, express the truth-conditions of their utterances, and are as a result unevaluable with respect to a possible world. These sentences do not, at least at first sight, display any form of ambiguity, vagueness, context-sensitivity or syntactic ellipsis; on the other hand, the idea that their utterances really express a determinate content is very much implausible. These were the sentences which I characterised as semantically under-determined in both the **Introduction** and in **Chapter 1**. It is very important, in order to demarcate the field, that we concentrate on exactly *this* group of cases.

In particular, a sharp distinction should be maintained between them and the following, very different cases:

Non-sentential assertions. It has been argued (see Stainton & Elugardo 2004, Stainton 2005) that genuine assertions can be performed even employing non-sentential expressions. Uttering single phrases or even words can qualify as the performance of an assertion. Suppose a speaker, say a coffee shop attendant, utters (1) while showing a coffee pack to a customer; or suppose a child utters (2) while handing an empty glass to her mother.

(1) From the mountains of Colombia
(2) Water

In these cases, even though they are not using sentences, speakers have been taken as performing genuine, truth-evaluable assertions. Since there is not enough syntactic structure in order to allow a semantic interpretation of (1) and (2), the role of pragmatic enrichment in determining the content of these assertions has been pointed at as crucial. As Stainton & Elugardo write

> context must be contributing to the truth-conditions of the assertion in ways that are not traceable to items in the disambiguated syntax, i.e., to slot-filling […] . In which case, […] free pragmatic enrichment must be playing a role in determining the literal truth-conditional content of [the] assertion. (Stainton & Elugardo 2004: 445)

The difference between the cases at interest for me and the ones focussed on by Stainton & Elugardo is straightforward: the former are examples of sentential constructions, while the latter are examples of non-sentential constructions. From the point of view of a proponent of free pragmatic enrichment, the difference is not particularly significant, in that pragmatic mechanisms of enrichment clearly operate irrespectively of whether the linguistic object at issue is a well-formed sentence or not. From my point of view, the difference is crucial. Take the phrase "From the mountains of Colombia": here it seems that, prior to exercising one's semantic/conceptual competence, one has to reconstruct the (presumed) syntactic structure of the expression. This, however, is a process that involves investigating the speaker's intentions first. It is only after having settled these aspects that conceptual considerations may apply. So in the case of (1), conceptual constraints are not immediately of help because the priority goes to processes that are purely pragmatic, in that they aim at discovering facts that depend on the speaker's intentions, in order to assign a (even minimal) syntactic structure to the expression. Things get even worse with an utterance of (2). Here further structural aspects need to be settled, among which whether (2) is an assertion or any other kind of speech act (a question, a command). So once again it seems that questions as to the speaker's intentions have to be established *prior* to any application of conceptual constraints, and that conceptual constraints are not immediately of help here.

Be that as it may, the question of whether there are non-sentential assertions, and of which mechanisms help determine their content, is not our focus here (for views that contrast with Stainton's see Stanley 2000, but also more classic authors such as Dummett (1973: 297-8) and Evans (1982)); it will suffice to say that the conceptual constraints proposal is *not* meant to provide an (immediate) answer to these issues.

Evaluable, but Conversationally Inappropriate Assertions. Another category which should not be conflated with the one I am presently dealing with is the one comprising sentences that express full-fledged propositions, even though these propositions are somehow inappropriate from a pragmatic or conversational point of view. Below are some examples:

(3) It will take time to get there;
(4) She insulted him and he left;
(5) Sam has a good writing;

An utterance of (3) states that it will take a certain amount of time to get to the demonstrated place, which is rather trivial; the content of the utterance is fully evaluable (indeed, true), however it is of course not what speakers may really intend their hearers to understand, the intended content being that it will take *more time than expected* to get there. The literally expressed content is, therefore, fully propositional (and evaluable), but uninformative because trivially true, and hence inappropriate from a strictly conversational point of view. Something similar happens with (4). The content expressed by the sentence is that a demonstrated woman insulted a demonstrated man and that the man left, but it does not specify either the order or the relation between the two events. This is rather odd, because plausibly a speaker who utters (4) intends the hearer to understand that the leaving happened after and it was caused by the insulting. Although the content literally expressed by (4) is evaluable (it has the truth-conditions of a logical conjunction), it may therefore be conversationally inappropriate. Example (5) also belongs to the category of evaluable, but conversationally inappropriate assertions, if placed in the right context: take and utterance of (5) expressed as the content of a recommendation letter for a candidate Ph.D. student: though fully evaluable, it goes off-topic in that it doesn't engage in an evaluation of the skills typically distinctive of an academic fellow. The fault of any utterance of (5) in a context like this is that of not being relevant, or of giving too little information, and hence of being conversationally inappropriate.

The kind of under-determination exhibited by sentences such as (3)-(5) is not, as I have argued both in the **Introduction** and in **Chapter 2** (**section 5**), adequately characterisable in terms of "deep" semantic under-determinacy. "Deep" semantic under-determinacy obtains just in case the content that a sentence expresses, besides not fully specifying what the speaker intends, does not specify a truth-condition *at all*. According to this characterisation, cases such as (3)-(5) are not cases of "deep" semantic under-determinacy. They are instead cases of "shallow" semantic under-determinacy, in which a sentence's meaning merely fails to express what the speaker intends, but does express an evaluable content anyway.

As one may have noted, I have put together examples which, according to some theorists, belong to the category of "explicature", like (3)-(4), with others that may be viewed as more germane to Gricean implicatures, like (5). There may be differences between the two sets of examples[42], however for present purposes I shall focus on their similarities, which are worthy enough of attention. When confronted with an utterance of any of (3)-(5), the hearer has to engage in a form of inference which plausibly takes into account, together with the literally asserted content, also information concerning the conversation and its (tacit) rules, the speaker's intentions and all those elements that may be salient in the context, in order to identify the message really conveyed over and beyond the sentence uttered and its literal meaning. This happens to an equal extent in the case of explicatures and implicatures, as more than one author admits (see Carston 2002a). The only difference that is invoked is between an implicit, spontaneous, unconscious inference in the case of explicature and a more explicit, effortful and relatively slow reasoning in the case of implicature. These considerations all suggest that an explanation of how the content conveyed by utterances of these sentences is arrived at by hearers doesn't rely primarily on considerations connected with *conceptual constraints*. Here what seems central to content individuation is finding out what the speaker intends to communicate; conceptual considerations may apply, but only as a consequence of that. Again, then, I shall say that the conceptual constraints view is not intended to provide a

---

[42] According to Carston (1988), implicatures typically do not imply explicatures, i.e. they cannot be logically "stronger" than them; this would make explicatures cognitively redundant. According to Recanati (1989), implicatures do not fall within the scope of logical operators, while explicatures (or "what is said") do. Recanati (2004) also emphasises the fact that implicatures are derived from a well-defined sentence-content, while "what is said" is typically the product of the enrichment of a sub-propositional sentential content.

treatment of examples such as (3)-(5), which are better dealt with by means of an account in terms of pragmatic processes which proceed through intention-ascription.

How large is, then, the scope of the account I am defending? Are there any other examples of semantic under-determinacy besides the ones I have been referring to all throughout this work? I believe there are countless examples of sentences analogous to those considered so far—indeed I provided a good number of them in the **Introduction**. There seems no better way of characterising semantically under-determined sentences as those sentences which do not *prima facie* look as if they suffer from ambiguity, vagueness, indexicality or ellipsis, and about which it seems wildly implausible to suppose that they really express a determinate, invariant content across contexts. According to Bach (1994a), these sentences are set apart from others in that their content is, as it were, *conceptually truncated*, i.e. in that some essential aspects of the concepts involved in the expressed content are left unspoken, or unarticulated. Consider for example:

(6) Olivia is tall;
(7) Paracetamol is better;
(8) I'll see you outside;
(9) The guests are late.

In all these cases, the content expressed by the sentence is, as it were, conceptually incomplete. Being tall is being so with respect to a *comparison class*; being better is being such *in contrast to something else*; being outside is always *relative to some interior location*; while being late is invariably being so with relation to a certain *schedule*. All these relational aspects are left out by the compositionally derived semantic contents of (6)-(9): that is why Bach attributes to the content of these sentences an incompleteness which is conceptual in kind. My view is partially in agreement with Bach's insight, for I do believe that semantic under-determinacy could be depicted—at some level of analysis—as conceptual truncation. My account though takes a further step in that it speculates that the very process of comprehension of (6)-(9) and the like is accomplished thanks to the hearer's exploiting of her underlying competence about the concepts (meanings) involved, and not thanks to the hearer's employment of Gricean, fully pragmatic mechanisms.
To sum up, in this section I have circumscribed the range of phenomena for which I am going to propose a treatment in terms of conceptual constraints. Firstly, I have ruled out so called *non-sentential assertions*, in that I believe that conceptual considerations are secondary with respect to purely pragmatic considerations which may be used to establish a (albeit minimal) syntactic structure for the expressions involved. Secondly, I have ruled out all cases of *"shallow" semantic under-determinacy*, i.e. those cases in which a sentence expresses an evaluable semantic content which is however pragmatically inadequate. I have admitted as amenable to a conceptual constraints treatment all those cases in which we have a sort of semantic under-specification which doesn't seem to be imputable to any clear case of ambiguity, vagueness, indexicality or ellipsis; these cases seem to involve "conceptual truncation"—something which suggests that a treatment in terms of conceptual constraints might be on the right track.

2. Some Background on Utterance Understanding
After having delimited the scope of application of the conceptual constraints account, let me compare the way the conceptual constraints view approaches utterance understanding with other approaches to the same topic that may be considered similar to it. The aim is, once again, that of providing more details concerning the proposal I am currently advancing. In this section, I shall sketch the main positions in the debate concerning what kinds of processes make it possible to achieve an understanding of utterances of under-determined sentences in context. The point of departure for all these approaches is that linguistic meaning under-determines speaker's meaning (or

what a speaker means with an utterance) in some "deep" way. It is up to the theorist to indicate the —linguistic or extralinguistic—means thanks to which communication often turns out as successful, notwithstanding the lack of determinacy on the part of language. The main approaches that have been defended are essentially three: inferential approaches; non-inferential approaches and syntactic approaches.

2.1 Inferential accounts

The first family of accounts identifies *inferential processes* as the *trait d'union* between linguistic meaning (i.e. the content of a sentence compositionally determined) and speaker's meaning (i.e. the content of an utterance of that sentence in the context in which the utterance is performed). The main claim of an inferential account may be identified with the following:

[Inferential account]
What allows a hearer to derive the speaker's meaning *m* of a certain utterance *u* of a sentence *s* is invariably an inferential process.

The speaker's meaning of an utterance is not simply the sentence's content-in-context: even though the latter may incorporate elements that depend on the speaker's intentions (like for instance the referents of demonstrative expressions such as "that", "there"), it is usually not considered an example of speaker's meaning. Speaker's meaning requires that the content that the utterer intends to convey departs to a greater extent from the sentence's conventional content; incorporating the intended references of some expressions is not enough. For this reason, typical examples of speaker's meaning are metaphors, metonimies, implied meanings (Gricean implicatures), indirect speech acts; less obvious examples of speaker's meaning are provided, for example, by the use of quantifier phrases, definite descriptions, relational terms ("tall", "better"). Among these expressions one may list, interestingly for our purposes, the semantically under-determined expressions that have been surveyed. I take it as a distinctive assumption of an inferential account that, save for clear cases of indexicality, the "gap" between sentence meaning and speaker meaning is filled in inferentially for *all*, more or less typical, cases of speaker's meaning.

2.1.1 The Gricean approach

Grice's seminal work *Logic and Conversation* may be taken as the point of departure of an inferentialist approach to utterance comprehension. Grice's main concern in the article is to account for the "divergencies in meaning between […] formal devices […] and […] their analogues or counterparts in natural language" (22). His explanation of how this divergence is overcome relies on the distinction between sentence meaning and speaker meaning: the former is the semantic content of a sentence, determined on the basis of convention; the latter is the content that the speaker intends to convey with her utterance of that sentence. While linguistic competence (*plus* knowledge of the context) is sufficient for a hearer in order to figure out sentence meaning (in context), it doesn't suffice in order to figure out speaker's meaning. Further competences need to be involved, which are not linguistic in nature: these are distinctively inferential competences, together with other capacities that we may call "meta-representational". Let us introduce Grice's model of utterance comprehension with an example. Suppose John asks Mary whether she would like to go to the cinema. Mary replies "I am tired"; the sentence's conventional meaning doesn't allow John to fully grasp what Mary means. On the assumption that Mary is being cooperative, i.e. that she is striving to give her best possible contribution to the conversation (see **Chapter 1**), John could reason as follows:

1. Mary said that she is tired;
2. She wouldn't have said that she is tired unless she believed that she doesn't want to go to the cinema;

3. She knows that I can suppose that she might believe that she doesn't want to go to the cinema;
4. She hasn't done nothing to stop me believing that she believes that she doesn't want to go to the cinema;
5. Therefore, she intends me to believe that she believes that she doesn't want to go to the cinema;
6. Therefore (by saying that she is tired) she means/implies that she doesn't want to go to the cinema;

The first important feature of Grice's "working out" schema is that it is an instance of *inference to the best explanation*, where the *explanandum* is Mary's conversational move as described in premise (1). The assumption that is operative in the background here is the presumption of speaker's *cooperativeness*; that is, the performance of the whole inference has as an implicit condition that the speaker is striving to give the best possible contribution to the conversation, even though appearances may suggest otherwise. Thus, Mary's response may appear to be non-cooperative, in that it violates some of the conversational maxims, namely the ones that recommend relevance and clarity. The hearer assumes that, despite this apparent violation, the speaker is nevertheless cooperating to the communicative enterprise; with this assumption in the background, the hearer follows a train of reasoning that leads him to conclude, as the best explanation for (1), that the speaker is implying that she doesn't want to go to the cinema.

The second important feature of a Gricean inference like (1)-(6) is that its premises and conclusion all feature an *attribution of attitude* to the speaker, as e.g. a belief-attribution or an intention-attribution. Attributing a belief, or intention, or desire to another subject is a way of representing what is going on in that subject's mind, i.e. it is a way of representing another subject's mental representation. Let us call any representation of another's mental representation a *meta-representation*. Thus, the inferential schema set forth by Grice contains, at each step, a meta-representation, such as "Mary believes that φ", "Mary means/ implies that φ". Humans are naturally capable to form and entertain meta-representations, so Grice's model is plausible insofar as any fully developed human cognitively equipped so as to mentally form and entertain an attribution of speech, belief and intention to another subject can be thought of as going through one of these inferences.

Several theorists have criticised the Gricean model in the account that it might be regarded as *cognitively too demanding*, and therefore as psychologically implausible. Sperber (2000), Wilson (2000) and Sperber and Wilson (2002) forcefully highlight this limit. In the following excerpt, for instance, Sperber rhetorically stresses the

> […] cumbersome character of the inferences that Gricean pragmatics necessitates every time a speaker's meaning diverges from sentence's meaning [...]. Do we really have, in the case of implicature, of indirect speech acts, of metaphor, or of irony, to reflect on what the speaker knows we know she knows, on what respecting the maxims requires of her, on what she might mean and not mean? Does it take more effort and hence longer to process such utterances? [...] Do we want to attribute to young children these complex inference patterns or to deny them the ability to comprehend metaphor and other forms of so-called indirect speech? As a rational reconstruction of how inferential comprehension might be possible, Grice's account, though not without problems, is certainly appealing. As a psychologically realistic account of the mental processes actually involved in comprehension, it is much less so. (Sperber 2000: 131-2)

Critics of Grice argue that, as far as the use of *inference* is concerned, it is implausible that even fully developed adults ordinarily interpret utterances by performing such complex inferences as the one in (1)-(6), in an overt and in a step-by-step fashion. That is, it is implausible to suppose that listeners explicitly formulate premises such as (1)-(5) in their mind, explicitly follow a linear train of reasoning and explicitly reach the conclusion at (6).

The same considerations have been advanced with respect to *meta-representation*: critics complain that it is implausible that fully developed adults consciously or explicitly formulate the meta-representational propositions that feature in (1)-(6), for this would affect negatively the speed and easiness of the performance; quite to the contrary, it seems that if we entertain any meta-representations, we do that implicitly or "by default". Significant evidence on this point comes from developmental psychology. Experiments have shown that pre-verbal infants, who most certainly lack the capacity to explicitly form meta-representations, are nevertheless capable of inferential communication (see e.g. Bretherton 1991). This suggests that meta-representations are taken care of by capacities that are implicit, covert and probably *much less sophisticated* than the ones that are involved at a conscious level in that, for example, they do not presuppose either the capacity to reason consciously or any linguistic capacities.

As a reaction to the Gricean model, some theorists have started considering a view that countenances utterance comprehension as an inferential process which takes place at a *sub-personal, implicit, unconscious* level, and which is characteristically fast, painless and automatic. This view usually goes hand in hand with a Fodor-style assumption as to the architecture of the mind, which distinguishes between *modular* components and *central system* components (see Fodor 1983). Modules run their activity below the level of consciousness, they respond to and manipulate only certain specific inputs; moreover, they characteristically work in isolation, i.e. they do not communicate with each other and are "blind" to any stimulus that lies outside their domain. On the other hand, the central system has access to the outputs (but not to the internal processes) of each single module and typically integrates information coming from different domains. Its activity may be run at a conscious, reflective level, and it is responsible for what we may identify as reflective, rational thought. On the assumption that Gricean inferences may be implemented by central system mechanisms, this resulting in slow, effort-taking and too sophisticated instances of reasoning, some theorists oppose the Gricean model by arguing that utterance comprehension processes take place at a "modular" level, which explains their being unconscious, automatic, fast and non-effort-taking. In the next section, I shall present the most notorious example of anti-Gricean, "modular" approach, known also as Relevance Theory.

2.1.2. The Relevance-Theoretic Approach

According to Relevance Theory, communication through language consists in essentially two activities: the first one is an encoding-decoding activity, and the second one is an inferential activity. The encoding-decoding task is what humans undertake when they produce utterances of meaningful sentences (or pieces of sentences) in a language. Language is a code to the extent that it consists of a collection of symbols provided with a conventional meaning; encoding a message means "putting into words" that message; de-coding a message means grasping the content of (or "thought" expressed by) a certain sequence of signs by computing together their meanings. The encoding-decoding activity usually leads to rather partial results: often, what is encoded in a certain occurrence of a sentence is only part of what the speaker means by uttering that sentence[43]. For example, if A and B are both in a shopping mall, and A utters "John is outside", an important part of what A means with her utterance is not encoded in what A said—namely that John is outside of the shopping mall A and B are both in. How can the addressee work out the content of A's utterance?

Here the *inferential* part of the communication process emerges as key, for inference is thought of as being sufficient for filling the gap between the defective sentence meaning and the intended content of the utterance. Though the Relevance-theoretic account shares with Grice a strong emphasis on inference, the crucial point of departure from Grice is that, for Relevance theorists, the

---

[43]   It is a matter of debate whether sentence meaning *needs not* encoding all the speaker's meaning as opposed to whether sentence meaning *cannot* encode it. If one thinks that it simply needs not, one is subscribing to a view of semantic under-determinacy as merely *convenient* (see e.g. Quine, 1960: 193-4, Katz 1972); if one believes that linguistic meaning can't encode speaker's meaning, one is endorsing a view of under-determinacy as *essential* to language (see Carston 2002).

inferences at issue cannot be explicit and reflexive, but are conceived as *implicit or subpersonal,* that is they go on "unconsciously, painlessly, and fast" (Sperber, 1995). The difference between a Relevance-theoretic inference and a Gricean one is, as mentioned previously, equivalent to the difference between a process carried out by a Fodorian "module" and a process carried out by a "central system".

In Relevance Theory, the inferential processes which underlie utterance comprehension are characteristically aimed at obtaining a *relevant* result. Relevance is defined as the property carried by a stimulus whenever the cognitive processing of that stimulus has the highest possible cognitive effects at the lowest computational cost (see Sperber&Wilson 1986/1995: 153). The Communicative Principle of Relevance states that every communicative act (be it an ostensive gesture or the articulation of a sentence) carries with it a presumption of relevance, i.e. it by itself communicates to the addressee that it is worth for the addressee to attend to it. Given this presumption of relevance, the hearer is entitled to process the content of the communicative act in the following way: by following a "least effort" method, she identifies the interpretation which best satisfies the expectation of relevance raised by the input stimulus. Once this interpretation is reached, she stops.

A qualification here is in order: Even though they put much weight on *implicit* and *unconscious* trains of reasoning, Relevance theorists are certainly committed to a standard notion of what an inference is. According to Sperber (1995), inference "consists in starting from some initial assumptions and in arriving through a series of steps at some conclusion". To go back to the example, we may depict the hearer of A's utterance of "John is outside" as going through the following piece of reasoning (here I follow Sperber 1994): let us suppose that the first assumption consists in the proposition that A said "John is outside"; the second assumption is to the effect that both A and B are in a shopping mall; the third assumption is to the effect that what the speaker said is maximally relevant. From these assumptions, the hearer may, through a certain number of steps and following the path of least effort, arrive at the conclusion that what A means is that John is outside of the shopping mall. This all being said, it should be added that, even if Relevance-theorists are, at an abstract level, committed to a standard notion of inference, it is also largely accepted that, in practice, the inferences we use for utterance comprehension may not be a straight progression from premises to conclusions. As Carston makes clear:

> the process need not progress strictly logically from the accessing of premises to the drawing of conclusions. For instance, a particular conclusion, or type of conclusion, might be expected on the basis of considerations of relevance and, via a backwards inference process, premises constructed [...] which will make for a sound inference to the conclusion. The process may involve several backwards and forwards adjustments of content before an equilibrium is achieved [...]. (Carston 2002a: 139)

A further characteristic element of the Relevance-Theoretic approach lies in that the inferences employed in utterance understanding are performed by a component which is a sub-part of the human capacity to *read others' minds*. As proponents of the Relevance-theoretic approach hold, humans have developed a habit to interpret the behaviour of other subjects in terms of what they *believe or intend*; that is, they have developed an ability to explain others' behaviour by forming and entertaining *meta-representations*, i.e. representations of their fellows' experiences, thoughts, desires, intentions. Meta-representation is used also in the interpretation of communicative behaviour, which is thought of as a sub-species of intention-driven behaviour. When the interpretation of linguistic uses is at issue, Relevance theorists hold that what is usually called "pragmatic processes" of interpretation are inherently meta-representational, that is, they aim at working out *what is meant* or *intended* by the speaker. As Carston writes:

> [T]he pragmatic comprehension system interprets communicative behaviour in terms
> of an intention on the part of the speaker to bring about a certain belief state in the
> addressee. Currently, the idea is being developed that [...] the relevance-based
> comprehension module may be a sub-module of the more general mental-state
> attributing module [...]. (Carston 2002a: 132)

Here the crucial contrast with the Gricean account is owed to the fact that meta-representational
pragmatic processes are assumed to be covert or "modular", while Grice's "working out" inference
may as well be a piece of overt, reflective reasoning. With respect to meta-representation, the two
models differ in the *level of sophistication*: a Gricean inference has both meta-representational
premises and conclusion, thus requiring the performer of the inference to have quite developed
representational capacities. By contrast Relevance-theoretic inference, being carried out by a
modular component, only needs to give a meta-representational conclusion (*X means that φ*), while
its premises need not be meta-representational—in fact, they do not even need be accessible to one's
awareness. As Sperber writes:

> The conclusion of [...] an interpretation process is an attribution of a meaning to the
> speaker, hence a metarepresentation, but the premises in the inference process need
> not be metarepresentational. Therefore this procedure can be followed by a relatively
> unsophisticated metarepresenter. (Sperber 2000: 133)

The scientific hypothesis that underlies Sperber's claim is the following: at some point in their
evolution, humans have become endowed with a "mind-reading" or "intention-detecting" module,
which is responsible for the capacity to explain other subjects' behaviour in terms of beliefs, desires
and intentions. As communicative practices arose, this module evolved into an even more specific
"pragmatics" sub-module, devoted to the detection of communicative intentions. The core thought
is that, the component hosting this "pragmatics-related" mind-reading capacity being a module, the
inferential processes that it carries out are typically input-specific, intuitive, spontaneous,
subpersonal and *not even explicitly meta-representational*. This starkly contrasts with the Gricean
picture, which depicts pragmatic processes of utterance interpretation as depending on the capacity
to *explicitly* form and reason with meta-representations. Schematically, the difference between
Relevance-theoretic inferences and Gricean inferences could be thus represented:

Gricean Inference	**Relevance-Theoretic Inference**
Stimulus: utterance of "ψ"	Stimulus: utterance of "ψ"
	Pragmatics Sub-Module
Input/Premise: *X said that ψ*	Input/Premise: *ψ*
.	.
[Reflective reasoning]	[Path of least effort]
.	.
Output/Conclusion: *X means that φ*	Output/Conclusion: *X means that φ*

To sum up, in this section I have surveyed the so-called "inferential accounts" of utterance
understanding: these accounts explain the success of verbal communication with the idea that
inference compensates the under-determinacy that is proper to language. Two approaches to an
inferential account have been identified: the Gricean approach and the Relevance-theoretic
approach. The Gricean approach regards inference in utterance understanding as a conscious,
explicitly meta-representational process; by contrast, Relevance-theory postulates the existence of a

"pragmatics-dedicated" module which terminates with a meta-representational conclusion (*X means that φ*) after having run sub-personal and not overtly meta-representational inferential steps.

2.2. Non-Inferential Accounts

A position which strongly contrasts with the account just presented has it that understanding tokens of sentences of a natural language is an altogether *non-inferential* activity. According to Ruth Millikan, for example, there is no inferential step between hearing an utterance and knowing what state of affairs it speaks about. Language understanding is as direct as perception. Of course, the reflective speaker could voluntarily focus on the items that serve as a vehicle—the phonemes or graphemes, the words, their meanings, the sentence's syntax—and from those infer the content of the uttered sentence. However, this is likely to happen only in a minority of cases, and it implies performing a less natural or spontaneous task. As Millikan states,

> you can directly perceive the phonemes or directly perceive the words being uttered by a speaker if you want to. But usually you perceive only the world affairs spoken about. (Millikan 2004: 121).

Even though language understanding doesn't require inference on the part of the hearer, still this doesn't mean that there are no processes whatsoever that allow the hearer to go from the auditory perception of a certain utterance to the perception of the state of affairs spoken about. These processes are not inferential, but rather merely *translational*: they involve moving from a representation to another representation. In particular, these processes are not inferential in the sense that: (i) they cannot be modelled in terms of premisses and conclusion[44]; (ii) they involve "functionally insulated" signs, i.e. items which are designed to serve only one purpose (or very few ones) and can interact only with items very limited in number and format.

The translational model returns a picture in which language is a system of symbols that can be used without a user being fully aware of or focussed on its structure (children and even most adults, for example, do not know what phonemes are, nor they master a semantic vocabulary); hearers can employ this medium in order to have access, via translation, to the cognitions and perceptions of the speakers and, ultimately, to the states of affairs that they talk about:

> Believing what one hears is exactly like any other way of gaining knowledge by perception without inference. [...] Coming to know something by believing what someone else says is making use of another instrument that extends perception—[...] the [...] perceptual and cognitive system of another person, through which some portion of the world has been carefully focused and then projected again via a medium that one knows how to interpret. (Millikan 1984: 304-5).

In such a framework, language gives us access to others' cognitive and perceptual systems, thus literally *extending* our own perceptual system. This has as a consequence that there is no difference between seeing that P and being told that P. In both cases we have a *perceptual* access to the state of affairs that P. Millikan is very explicit on the matter when she writes:

> Rain does not sound the same when heard falling on the roof, on earth, on snow, and on the water, even though it may be directly perceived as rain through any of these media. Exactly similarly, rain has a different sound when the medium of transmission is the English language ("It's raining!"). And it sounds different again when the medium of transmission is French or German. (Millikan 2004: 122)

---

[44] Note that this is different from the Relevance-theoretic claim (see section 2.1.2) to the effect that inferences need not be a linear route from premises to conclusion, for this view still subscribes to a standard notion of inference.

There are other ways for processes of utterance understanding to be non-inferential, besides being translational in Millikan's privileged sense. Recanati (2002b) considers cases in which a sentence's schematic meaning is "fleshed out" so as to be made propositional; he observes that the meanings that are subject to these processes of "free enrichment" or "completion" are too abstract and schematic in order to be considered conceptual representations. He then surmises that these mechanisms of supplementation, since they manipulate non-conceptual (or non *fully* conceptual) contents, are not inferential. Suppose John utters "It's raining", meaning that it's raining in Paris. This sentence expresses the content *that it's raining*, which exhibits a conceptual gap— corresponding to the locational aspect. The mechanism of completion by which the content gets "fleshed out" is considered by Recanati as non-inferential to the extent that it takes as its starting point a content which is not (fully) conceptual.

A second set of processes he considers are what he calls "transfer". Semantic transfer is involved in non-literal talk, and it is typically non-truth-preserving. Consider for example the transfer involved in understanding an utterance of "The ham sandwich is getting restless": it brings the hearer from the false proposition *that the ham sandwich is getting restless* to the true proposition *that the ham-sandwich-orderer is getting restless*. In these cases what connects the first, semantically expressed, proposition to the second, intended one, is an "associative" link which is by no means truth-preserving, and that is surely not a logical entailment.

To sum up, in this section we reviewed three ways in which utterance understanding could be regarded as non-inferential: One could claim, like Millikan, that understanding an utterance involves merely *translational* processes, but also, like Recanati, that it involves *free-enriching* and *associative* processes. Free enriching processes count as non-inferential to the extent that they manipulate non (fully) conceptual material, while associative processes are not inferential insofar as they may not preserve truth.[45]

2.3. Syntactic accounts

An altogether distinct account of how utterance understanding works may be derived from Stanley's (2000) claim that all pragmatic effects are traceable to logical form. Let us consider an utterance of "It's raining"; as many theorists have written, the location of the rain, which may be part of the truth-conditions of the utterance, is not articulated in the syntax of the sentence. As we saw in **Chapter 3**, authors such as Stanley, Szabó and Rothschild & Segal argue that these elements are really articulated at a *covert* level: They have no phonetic or graphic realisation, however they are present in the deep syntax, or logical form of the sentence at issue; at least in contexts in which no variable-binding operators occur, they could be represented as *unfilled positions* or *free variables*, subject to either variable assignment in context or to binding by operators. Utterance understanding, then, becomes predominantly a matter of slot-filling, of "saturating" some pre-determined syntactic structures. Sentences are represented as partly filled syntactic templates, where empty argument-places are to be filled-in only with a restricted set of potential values, on the basis of what is relevant in context.

---

[45] The distinction between inferentialists and non-inferentialists is relevant also in the debate concerning the epistemology of acceptance of an interlocutor's assertion—or *testimony*. Non-reductionists about testimony claim that we are *prima facie* justified in believing what we are told, in absence of reasons to believe otherwise; for Burge (1993) and McDowell (1993), this implies that there is no inferential process that the hearer needs to go through in order to acquire a justification to believe a testimonial report. Reductionists, on the other hand, think that justification for believing what others say comes from the subject having independent, non-testimonial reasons to believe; these independent reasons may, for some theorists, come from *inductive inference* as to the reliability of the source (see Lyons 1997, Rysiew 2000), as well as from *inferences to the best explanation* (see Harman 1965, Fricker, 1995).

Proponents of the syntactic approach claim to be supported, in a relevant array of cases[46], by evidence provided by contemporary research in syntactic theory. Radically pragmatic explanations, such as those appealing to "pragmatic intrusions" or "free enrichment", are regarded as methodologically flawed in that they straightforwardly choose strong non-semantic, unorthodox claims, without previously making sure that there are no possible syntactico-semantic (and hence more orthodox) accommodations of the phenomena they are interested in. As King&Stanley claim:

> Suppose a theorist is in the position of arguing that the semantic content of a certain construction must be augmented by pragmatic information to account for natural readings of that construction. To make out this claim, it is incumbent upon the theorist to provide a sketch of the correct syntax and semantics of the relevant construction. For it is only then that one can evaluate the claim that the syntax and semantics for the relevant construction does not by itself deliver the reading at issue. Far too often, those who advocate the thesis that a certain reading of a construction is due to pragmatics rather than semantics fail to live up to their obligation, and so make claims that are highly speculative and hence difficult to evaluate. (King & Stanley 2005: 147)

As we have seen in **Chapter 3**, not all approaches in terms of syntactically articulated hidden positions may be immune to criticisms of a methodological kind; indeed in some cases, positing hidden syntactic elements may turn out as methodologically unwarranted, if not theoretically inconvenient (in that for instance it increases the complexity of the theory in an undue way). The debate concerns largely empirical questions, and settling the dispute ultimately requires working on a case-by-case basis. Be that as it may, in what follows (see **section 5**), my point against a syntactic approach to comprehension will build on the idea that, if no syntactic/empirical claims about the presence of hidden positions in logical form can be expected from arguments so far provided (especially, from Stanley's Binding Argument), this doesn't compel the opponent of the syntactic approach to embrace a completely unorthodox, pragmatic account. The reason is that, in the view I wish to propound, contextual effects may be semantically/conceptually constrained, rather than being driven by purely pragmatic considerations, since I argue that conceptual/constraints determine *semantic effects which are however free from syntactic constraints*.

3. Inferences, after all

The aim of this section is that of ruling out the idea that the comprehension of semantically under-determined sentences such as "It's raining", "Jill is ready", "The leaves are green" etc. is inherently a non-inferential process like the ones envisaged by Millikan or Recanati. Arguably, understanding utterances of these sentences is a genuine inferential undertaking.

The view that inferential processes may be involved in the comprehension of, e.g. "It's raining" is not accepted by all authors. For example, Recanati (2002b) emphasises that a process can be said "inferential" only if it takes as its input a conceptual content (see **section 2.2**). Some sentences express contents that are very "abstract" or "schematic": this may make these contents not enough conceptually structured to enter an inference:

> One of the primary pragmatic processes involved in the determination of what is said consists in "saturating" or "completing" a schematic meaning in order to yield a truth-evaluable representation. All the inferentialists rightly emphasise the abstract,

---

[46] For example, it has been argued that the syntax of sentences (both in the finite and non-finite mode) articulates temporal elements (see King (2003), Enc (1987), Ogihara (1996)); the syntax of comparative constructions and of conditionals has also been argued to contain hidden syntactic positions. In King and Stanley's (2005) view, positing these syntactic elements helps making sense of the otherwise vague and unorthodox claims that there are "pragmatic processes" affecting the interpretation of these constructions.

schematic character of the "linguistic meanings" which serve as input to primary pragmatic processes. That schematic, abstract character renders somewhat dubious the claim that those meanings [...] are themselves "conceptual representations". I would rather say that they are non-conceptual [...]. (Recanati 2002b: 122)

The conclusion would be that, since they involve non-conceptual contents, processes of enrichment of the content of certain sentences are non-inferential. In the case of "It's raining", Recanati would say that, since the content *that it's raining* is too abstract or schematic, it is non-conceptual, and hence there can be no account in which the enriched content *that it's raining in Paris* is the result of an inference. The same would apply to "Jill is ready": the content *that Jill is ready* only is a propositional schema, i.e. a highly abstract and skeletal object; this, in Recanati's view, is enough in order for the content at issue to count as non-conceptual. The hearer who enriches the content of "Jill is ready" as *Jill is ready to ski* cannot therefore be described as having carried out an inference. I do not agree with Recanati's point of view. The fact that a content is "abstract" or "schematic" may make it under-determined, i.e. incapable of expressing a truth-condition in a context of utterance, but it doesn't entail that it lacks a conceptual structure. In the case of "It's raining", it seems to me evident that the content expressed by the sentence is a fully conceptually structured content. This content may not express a full-blown proposition in the context of utterance of the sentence, but this doesn't prevent it from entering in an inferential relation with other propositions. As far as I can see, the following is a perfectly acceptable inference, in which the conceptual structure of the predicate "to rain" becomes fully apparent. The starting assumption is the following: if one knows what "to rain" means, one knows that, e.g., raining involves drops falling from the sky. This implies that one is in a position to use the predicate "to rain" in the following inference:

[1] It's raining;
[2] If it's raining, drops are falling from the sky;
[3] Drops are falling from the sky.

That the predicate "to rain" can take part to inferences such as the one above means that it has a conceptual content that can be articulated and can explain the inferences the predicate becomes involved in. That predicates like "to rain" express a conceptual content, even if they cannot make the content of the sentences they occur in fully propositional, is essential for our purposes, for this allows us to account for the cases of enrichment we are interested in terms of *inference*.
Let us illustrate how such processes may be addressed as inferential. Suppose a speaker utters "It's raining". By the lights of formal, compositional semantics alone, the expressed content is <*that it's raining*>. This content is plausibly the object of the enrichment. The hearer could then reason as follows: first, she could invoke as a premise of the reasoning the conceptual constraint to the effect that, if it's raining, it's raining somewhere; then she could go on acknowledging that location X is under discussion; this would lead her to the conclusion that it's raining in X. Again, suppose that the speaker utters "Jill is ready". By the lights of formal, compositional semantics alone, the content expressed is <*that Jill is ready*>; this content being unevaluable, the hearer could reason as follows: "If Jill is ready, she is ready for something" (thus applying a conceptual constraint); "Activity X is under discussion; so Jill is ready for X". Or, again, suppose the speaker utters "The leaves are green"; formal, compositional semantics delivers the content <*that the leaves are green*>. This content being poor, truth-conditionally speaking, the hearer may reason as follows: "If the leaves are green, they must be green in some respect" (thus applying a conceptual constraint); "Being green in way X is under-discussion; so the leaves are X-green".
It goes without saying that the inferences countenanced here need not surface at the conscious, personal level. Still, let me make clear that, even if hearers do not explicitly run these inferences,

these pieces of reasoning look like acceptable rational reconstructions of processes that might be gone through in an explicit manner by any rational, competent speaker. That these processes are at least virtually available to any rational, competent language user would be enough to make the point that processes of "enrichment" of the contents of "It's raining" *et similia* need not be conceived as non-inferential, as if the transition from the sentence's meaning to the expressed content and truth-conditions in context were a matter of "direct perception" (as in Millikan's framework) or of merely associative or transfer processes. (As to the question whether *empirical tests* could shed any light on the issue, let me postpone the discussion to the **Conclusion** of this thesis, where I shall sketch some of the possible empirical developments of the conceptual constraints proposal).

To sum up, the main aim of this section was to argue that comprehension of utterances of "It's raining" and the like follows an *inferential*, rather than a non-inferential path. I first considered Recanati's criterion to the effect that a process is inferential only if it involves conceptually structured material. Then I argued that it is wrong to regard the "schematic contents" of certain semantically under-determined sentences as not conceptually structured. For example, "It's raining" expresses a very schematic content, but seems to me to exhibit a full-fledged conceptual structure; this is made apparent by the fact that sentences expressing these "skeletal" contents can nevertheless be used to run semantic inferences. Under-determination of content owed to the sentence's expressing a mere propositional "template" should not be mistaken as absence of conceptual structure. These considerations favour the view that understanding an utterance of "It's raining" is, after all, an inferential process.

3.1. Ruling out Conventional and Generalised Implicatures

Let me now move to an important qualification that I believe anyone who defends the position I am currently taking owes to the potential critic. One may point out that what I am dealing with, under the name of "enrichment through conceptual constraints" is either a *conventional* or a *generalised implicature*. Let me consider the idea of conventional implicature first. From Grice (1967/1989)[47], we know that a conventional implicature is part of what a speaker suggests (without saying it) with his use of a certain expression, for the simple fact that she has used that expression. This happens with "therefore": using this word in the sentence "He is an Englishman, he is therefore brave" raises the implicature that the man's being brave is a consequence of his being an Englishman; though this is not strictly speaking said by the sentence, use of the word "therefore" is sufficient to suggest that a relation of consequence obtains.

This being settled, can we say that, with the enrichment presently at issue, we are facing just another case of *conventional implicature*? One feature, it seems to me, sets the present case apart from cases of recognised conventional implicature. As Grice and others admit, the consequence relation suggested by a use of "therefore" is not part of the truth-conditions of the sentence "He is an Englishman, he is therefore brave"; it is simply a suggested aspect. If the person referred to by "he" were brave, but not as a consequence of being an Englishman, the speaker would still have said something true, although she would have *implied* something false. Now consider an utterance of "It rains"; even though the sentence may be true even if it's not actually raining in any specific place (as in "Whenever drops are falling from the sky, it rains" : see **Chapter 3** for what I called "zero readings") any felicitous use of such a sentence nevertheless evokes a locational component. Suppose one were to utter the sentence "Whenever drops are falling from the sky, it rains". Even though this is true even if it's not raining anywhere, still the fact that rain episodes are spatially located is a necessary truth. Suppose this necessary truth was somehow conventionally implicated by an utterance of "Whenever drops are falling from the sky, it rains". Being an implicature, it should be cancellable. However, cancelling this implicature results in plain incoherence. One cannot consistently claim "Whenever drops are falling from the sky, it rains, but not in any place". If the

---

[47] But see also the following references from introductory articles and books on the matter: Karttunen & Peters 1979; Levinson 1983: 127; Horn 2003: 383.

locational aspect were implicated, the sentence wouldn't sound contradictory, for implicatures (even conventional ones) have no effect on the truth-conditions of the uttered sentences. From this comparison I conclude that the alleged conventional implicature incorporated into the meaning of "It rains" is more than just an implicature: it is an essential aspect of the utterance's truth-conditions, such that if it were not to be included in them, the speaker would be culpable of having said (not just suggested) something plainly false.

There is a second possibility for the theorist who wishes to assimilate the present brand of enrichment to some already familiar phenomenon: saying that we are dealing with just another example of *generalised conversational implicature*. What is a generalised conversational implicature? Once again, we have to make reference to Grice (1967/1989) (but see also Levinson 2000), who claims that a generalised conversational implicature is a conversational (hence non-conventional) implicature, which has become customary, or habitual. For example, it has become customary, when uttering a sentence like "I saw him with a woman, last night", to imply that the woman in question is neither the wife, nor the sister or mother of the man referred to with the pronoun "he". Another example is the use of "Some" to imply "Not all": when a speaker says "Some of the protests were shut down" she usually implicates, and hearers normally take her as implicating, that not all protests were shut down. The implicature is elicited by default, but it is always cancellable. For example, the speaker could say "I saw him with a woman, last night. In fact it was his sister." and "Some of the protests were shut down; indeed, all of them were".

Is enrichment through conceptual constraints a case of generalised conversational implicature? It is hard to deny that every use of "It's raining" or "Jill is ready" or "The leaves are green" carries with it some customary, or habitual aspect. As Carpintero observes,

> When a speaker utters a token *i* of 'is raining' (as in 'it is raining'), he takes for granted that more information is available in context about the place he is referring to than that it is salient when the token *i* is produced. But it is a conventional fact about the use of the present tense with verbs meaning located events such as 'to rain' that the place referred to is indicated in that way, as it were, by default, and speakers rely on their audiences sharing with them this knowledge. (Carpintero 2001: 110)

Here Carpintero rightly notes that uses of "It's raining" take for granted that some salient information is already available in the context of the utterance for the interpreting audience. This is for sure a matter of custom or habit (or of "convention", in Carpintero's words). But does this suffice for the case to be an example of generalised conversational implicature? Here again, there is a fundamental difference between generalised conversational implicatures and the cases of enrichment at issue, namely that the former are *cancellable*, while the latter are not. To repeat, it is perfectly fine to retract what one suggested with the use of "I saw him with a woman, last night" by saying "I saw him with a woman, last night. In fact it was his sister"; or to cancel what one implicated with his use of the phrase "Some protests" by saying "Some of the protests were shut down; indeed, all of them were". This freedom is lacked by speakers who use sentences like "It's raining". In normal circumstances, it would make no sense to say "It's raining, but there's no place I am talking about where it's actually raining"; or "Jill is ready, but I'm not saying that there's some activity for which she is ready"; or, again "The leaves are green, but there's no respect in which they are such". These would not be just instances of implicature erasing; they would be plain contradictions. Note that the same holds even for *zero-readings* of "It rains", since even these uses evoke a locational aspect which would be erroneous to cancel as if it was an implicature: as already noticed, it would be a straightforward incoherence to utter "Whenever drops fall from the sky it rains, but not in any place". To conclude, then, there is evidence that conceptual constraints do not give rise to either conventional implicatures or to generalised conversational implicatures.

4. Reconsidering the Role of Meta-skills

As I made clear in **section 1**, the range of examples I intend to account for is rather limited: besides the repeatedly used "It's raining", "Jill is ready" and "The leaves are green", I provided another short list in (6)-(9), which adds up to the list provided in the **Introduction**. Understanding an utterance of each of these sentences in a context requires, as many—myself included—acknowledge, a certain effort, aimed to "compensate" the gappy content of these sentences.

In the last section, I confronted the question whether this "compensating effort" could be described as an inference. The whole section was devoted to arguing for a positive answer to this question. The inferential process may be reconstructed as follows: first, upon listening to an utterance of, e.g. "Jill is ready", the hearer is presented with an item whose content is, by the lights if compositional semantics alone, *<Jill is ready>* (premise [1]). The hearer then exploits her own conceptual knowledge to add premise [2] to the effect that, for instance, being ready is being ready for some purpose. If the hearer has access to the particular context in which the utterance is performed and knows what the speaker is talking about, then a further premise may be added ([3]) having to do with which activity is salient in context, for instance skiing. From these three premises, the hearer can draw the conclusion to the effect that Jill is ready for skiing.

[1] Jill is ready
[2] Being ready is being ready for some purpose
[3] Skiing is the salient purpose
[4] Jill is ready for skiing

As one may have noted, the inference I have written above does not involve any ascription of intention ("X intends that *p*") or belief ("X believes that *p*"), nor does it involve any speech-report ("X said that *p*"). All these are examples of *meta-representation*, i.e. of representation of the content of other subjects' attitudes and speech acts: beliefs, intentions, desires, assertions, utterances and so on. The ability to meta-represent is usually taken to require "mind-reading capacities" (see Sperber (2000), Wilson (2000), Sperber and Wilson (2002)).

Now I contend that, if the inference in [1]-[4] is a good reconstruction of a process really followed by hearers, then carrying out such an inference does not, *by itself*, require the use of any meta-representations of the form "X believes/intends that *p*", "X said that *p*", etc. The interpretation of utterances of such sentences such as "Jill is ready", "It's raining" etc. is therefore, *by itself*, not a task that requires the use of a hearer's mind-reading capacities[48]. We could describe an inference like the one in [1]-[4] as the "mechanical" application on sentence meaning of pre-determined rules for organising thought and discourse *plus* some collateral information as to what is salient in context. The whole task may well be performed by exercising capacities which are independent of our ability to "read our fellows' minds".

That inferential processes like [1]-[4] may well be, *by themselves*, meta-representation-free does not imply that these inferences cannot *interact* with meta-representational processes. Indeed there may be a role for meta-representation in *providing adequate grounds for one or more premises in the inference*. To illustrate, consider the implicit reasoning that a hearer has to go through when interpreting an utterance of "It's raining". In certain contexts, interpreting an utterance of "It's raining" may require that one identifies what location the hearer is referring to, or talking about, or presupposing as relevant. In general, I shall take the task of figuring out what a speaker is referring to, what a speaker is talking about, what a speaker is presupposing and so on as a mind-reading task, and therefore as a *meta-representational* task. If this is so, then in certain specific cases, a

---

[48]   After all, one may know that skiing is salient in context for reasons independent on one's meta-representing the beliefs and intentions of the interlocutor. Not to mention the fact that the hearer may employ non-meta-representational strategies in order to establish what is salient in context. For an account of how hearers may gain knowledge of the salient aspects of context while not taking the (putative) intentions of the speaker into consideration, see Gauker (2008).

meta-representational process may have to be run in parallel with the inferential, conceptually constrained process. In general, I shall say that whenever a conceptually constrained interpretation has to take into account such factors as what the speaker is referring to, what the speaker is talking about, what the speaker is presupposing, etc. a meta-representational process will have to be followed in parallel with a conceptually constrained process. So suppose that, in the context of utterance at interest, the speaker intends to talk about the weather in Paris. On the view I propose, the hearer will run *two parallel processes*: the first process consists in the application of the conceptual constraint to the effect that to rain is to rain somewhere. This process may be said to "expand" the semantic/conceptual structure of the sentence, thus opening, as it were, a new conceptual/semantic slot to be filled in with an appropriate locational value. The second process is a process aimed at figuring out *what location the speaker is talking about*: this process should be regarded as inherently meta-representational. Let us now see *how these two processes interact*. My proposal is that we regard the meta-representational process as yielding a conclusion that provides *evidential support* for one of the premises in the "mechanical", conceptually constrained inference. This premise shall have to do with what is going on in the context of the conversation, namely with what is salient or relevant, and may be identified with premise [3] in the inference-schema just laid out. I propose that the relation between premise [3] (or the like) and the meta-representation as to what the speaker is talking about will be such that the former is *justified* by the latter: that is, that the speaker is talking about, e.g., Paris shall justify premise [3] to the effect that Paris is salient in context. The relation is shown in the figure below:

[1] *It's raining;*
[2] To rain is to rain somewhere;
[3] Paris is salient in context;   **Meta-representation**: the speaker is
[4] *It's raining in Paris*.   talking about (the weather in) Paris;

To sum up what I have proposed so far: the interpretation of utterances of semantically under-determined sentences is an inferential, conceptually constrained and, therefore, "mechanical" process which, *by itself*, doesn't require use of the interpreter's meta-representational skills. Nevertheless, the process might interact with meta-representations, namely in all those cases in which the hearer has to figure out what the speaker is referring to, talking about, or presupposing. In these cases, meta-representations need not directly take part to the interpretive inference, but only need provide *evidential support* for some of its premises.

The resulting model is one where the role of meta-capacities is significantly reconsidered. Recall that, in the Gricean "working out" schema, meta-representation had a pervasive role, for both the inference's premises and its conclusion were meta-representations (of the form "X said that φ", "X means that φ"). In the Relevance-theoretic framework, processes resulting in meta-representations were carried out by modular components. This implied that the reasoning's premises needn't be explicitly meta-representational; nevertheless, the inference's conclusion was definitely a meta-representational one. The model I have proposed differs from both Grice's and the Relevance-Theorists' models in that, by postulating "mechanical" processes that respond only to conceptual demands, it envisages inferences whose premises and conclusions may be all non-meta-representational. If this is so, then *the role of meta-representation is largely reconsidered*: meta-capacities are involved in settling aspects that pertain to, e.g. what the speaker is talking about, but the resulting meta-representations need not also be premises in the interpretative inference and nor, for that matter, the inference's conclusion. Meta-representational elements may rather *justify* some of the inference's premises. The differences between the three inferential accounts may be captured schematically as follows:

Grice	Relevance Theory	Conceptual Constraints
Meta-representational premisses; Meta-representational conclusion;	**Non**-meta-representational premisses; Meta-representational conclusion;	**Non**-meta-representational premisses; **Non**-meta-representational conclusion; Meta-representations justify premisses;

(Note that the "conceptual constraints" position is represented in the chart as having it that language comprehension starts from *non*-meta-representational premisses and ends with *non*-meta-representational conclusion, thus being *wholly* meta-representation-free. However, what is really crucial about the position and sets it apart from both the Gricean and the Relevance-Theoretic approaches is the fact that comprehension is thought of as giving rise to non-meta-representational *conclusions*. Thus, it is not incompatible with my view that some of the premises in the inference may indeed be meta-representational, as long as they give rise to a non-meta-representational conclusion[49].)

An advantage of the conceptual constraints account is that it explains particularly well the fact that speakers seem to be able to interpret utterances of semantically under-determined expressions by performing conceptually constrained tasks and meta-representation-oriented tasks *independently of one another*.

Suppose that, for some reason, exercise of one's meta-representational capacities were impeded. There are cases in which comprehension could be accomplished anyway in a context, with the sole help of one's use of conceptual constraints and other collateral contextual (but not strictly speaking meta-representational) information. For example, imagine a Z-lander who hears an utterance of "It's raining". If she masters the concept of "raining" well enough and she independently knows that when it rains, it rains in Z-land, she may not need to meta-represent what the speaker is talking about in order to get the content that it's raining in Z-land. Or, suppose that an occurrence of "It's raining" is processed by a subject who doesn't have a capacity to meta-represent; if the subject has sufficient conceptual competence, she may still derive the content that it's raining somewhere, or she could use contextual information that she has independently of any meta-representation of the speaker's intentions in order to get the content that it's raining in (say) Paris. Or, finally, suppose one simply overhears an utterance of "It's raining", but one doesn't know anything about the intentions of the speaker. If one is a competent user of the concept (predicate) "to rain", one will be able to infer the content that it's raining somewhere. All these tasks I take to be possibly accomplished with the help of one's own independently possessed conceptual (semantic) competence, which in these cases is exercised altogether independently on one's capacity to meta-represent the intentions of the speaker as to what is meant with a certain linguistic act.

But the opposite scenario could also obtain: suppose that one's mastery of the relevant concepts were impeached. Still, one could employ one's meta-representational capacities in order to get an idea of what the speaker is talking about, and maybe, from that, "work backwards" until one identifies the conceptual constraints that should be applied. For example, suppose one were to hear the locution "movie snow". This expression may make it difficult to the listener to apply her conceptual competences about either the word "snow", or "movie". Still, suppose the hearer can meta-represent that the speaker is talking about techniques for faking weather phenomena in films; then she can work back to a modified, compound concept which contains novel constraints, such as FAKE-SNOW-USED-IN-FILMS.

One of the motivations for the conceptual constraints account lies in that it seems implausible that utterance comprehension relies too pervasively on meta-representation. Grice's account, which envisages interpretive inferences that are meta-representational through and through, was indicated as cognitively implausible (see Sperber 2000, Wilson 2000, Sperber and Wilson 2002). However, even the Relevance-Theoretic account assigns meta-representation a massive role, for it countenances interpretive processes as being preformed by modules that are dedicated to mind-

---

[49]   Thanks to Annalisa Coliva for having made this apparent to me.

reading, and hence are regarded as invariably resulting in meta-representational conclusions. My point in putting forward a conceptual constraints view is exactly to question the idea that the result of our interpretations of utterances be invariably meta-representations of the form "X means that φ", "X implicates that φ". Making this point implies advancing the idea that *not all comprehension is meta-representation or "mind-reading"*. There may be a species of comprehension which eludes the exercise of one's capacities to "read other subjects' minds".

I believe the latter insight points at an important philosophical distinction, which theorists engaged in investigating linguistic comprehension should definitely keep in mind. Supposing that the conclusion of a comprehension process is invariably a meta-representation such as "Mary means that q" is, I take it, tantamount to conflate comprehension with what one may call *"hermeneutics"*. Hermeneutics may be defined as the task of interpreting an utterance with the aim of settling the question of what the speaker/writer means; it is typical of those situations in which a disagreement may arise as to the speaker/writer's real intentions (what she "means with that sentence", "intends to convey with this choice of words" etc). In other words, hermeneutics is an interpretive enterprise whose final goal is entertaining a meta-representation like *X means that φ*. By contrast, comprehension may be characterised as the task by which one arrives to entertain, on certain grounds, an object-level representation *φ*. No doubt, it might be the case that, more often than expected, hermeneutics is needed in order to achieve and justify comprehension. This however is not a reason for merging the two tasks, especially when one is trying to capture the multifarious mechanisms by which language users can be said to "understand" the content of their fellows' linguistic productions. A proposal in terms of conceptual constraints is meant to emphasise precisely that not all interpretation of utterances needs result in hermeneutic, i.e. meta-representational conclusions. Some processes are simply comprehension processes, eventuating in object-level conclusions (even if this doesn't rule out that meta-representational, hermeneutic considerations can *support* processes ending up in comprehension results).

To sum up, in this section I have put forward the view that inferential processes of content-enrichment through conceptual constraints are meta-representation-free. The conceptually constrained interpretation of an utterance of "It's raining" in a context may be accomplished by a hearer with no need to exercise her mind-reading capacities, hence with no need to reason in terms of belief or intention attributions, or speech-reports. Meta-representation-free inferences may interact with genuinely meta-representational processes, however the latter may be viewed as simply providing evidential support for one or more premises of the former. The whole process and, most importantly, its conclusions, is conceivable as entirely meta-representation-free.

5. Against Syntactic Constraints: Standard cognitive formatting

In **Chapter 2**, the argument by means of which I drove attention onto the idea of conceptual constraints was structured as follows: Despite what Radical Contextualists say, we have an understanding of semantically under-determined sentences like "It's raining" or "Jill is ready", even out of a context. I called this kind of understanding *intuitive understanding*, i.e. knowledge of the truth-conditions of any utterance of the sentence, in any context. In addition, one could also achieve a "contextual" understanding, i.e. understanding of the intended truth-conditions of a particular utterance in a particular context. What makes intuitive understanding possible is, I suggested, our mastery of the relevant concepts invoked by the utterances—e.g. the concept of RAIN or READINESS (of course, as I already made clear, conceptual competence coincides with semantic competence for all those subjects who are competent language-users). So-called *conceptual constraints*, to be represented as ways of structuring thought and discourse, help us achieve an intuitive understanding of these sentences and may be regarded as also guiding our understanding of particular utterances in particular contexts.

The conceptual constraints hypothesis seems to be compatible with a syntactic approach to the comprehension of semantically under-determined sentences. The proponent of the syntactic

approach could take the remarks laid down in **Chapter 2** as suggesting that the sentences at interest really exhibit a more complex syntax than meets the eye. Further arguments may be appealed to in order to make the point stronger, such as Stanley's "Binding Argument". As I argued in **Chapter 3**, there are reasons to doubt the cogency of the arguments invoked by proponents of the syntactic approach; their commitments as to covert positions at the syntactic level arguably follow from a dubious methodology for syntacticians (see Neale 2007, Pupa and Troseth 2011, Collins 2007). By contrast, an approach in terms of conceptual constraints doesn't make any dubious commitments as to covert elements articulated at the syntactic level. The fact that the predicates "to rain", "being ready", etc. respond to constraints such as "to rain is to rain in some location", "to be ready is to be ready for some purpose", etc. is a *purely conceptual/semantic* fact, a fact that doesn't either suggest or reveal anything at the syntactic level. The application of conceptual constraints in figuring out the content of an utterance of "It's raining", "Jill is ready", etc. is not driven by syntactic demands, but by the need to reconstruct a content whose structure complies with *the way thought and discourse about the relevant matters is standardly organised*.

What does motivate the idea of such a standard cognitive/discursive "format"? Here is a story that could help one to gain a firmer grip on the idea. People belonging to the same social, cultural, linguistic community share a certain amount of information about their environment. The way this information is organised affects the way thought and discourse about numerous matters is itself organised. Yet, if every member of the community had her own way of organising thought and discourse, it would be extremely difficult to efficiently exchange information and coordinate action. In order for efficient exchange of information and inter-subjective coordination to be possible in general then, it seems that the most convenient route for members of the community is to develop *standard ways to cognitively organise thought and discourse* on a number of matters. The strategy could be equated with the adoption of a common currency which facilitates financial exchanges and trading. The process doesn't need to be the result of a collective decision or stipulation—it may ensue from adaptation. With this picture in the background, we may look at communication as a domain in which speakers utter sentences and hearers interpret them—also—by trying to relate the content of these utterances with some *familiar, standard format for thought and discourse*. (Note that this task is entirely different from—but not entirely incompatible with—a task of mind-reading, in that it is guided by pre-determined, pre-existing constraints and depends only marginally on the hearer's recognition of speaker's intentions).

The fact that a standard cognitive/discursive "format" characterises thought and talk about certain matters has an important implication as far as the relation is concerned between the words and the concepts that we use: at least when certain predicates are at focus, the view is committed to the claim that *the structure of the concepts (meanings) associated with the predicates is richer than the syntactic structure of the predicates themselves*. For instance, the structure of the concept RAIN is richer than the syntactic structure of the predicate "to rain".

One, however, may be skeptic with respect to the distinction suggested here between the syntactic structure of a predicate and that of a concept. How is it possible—one might wonder—that a certain concept has a more complex structure than that given by the argument places of the predicate used to express that concept? Isn't there supposed to be some sort of coordination or "harmony" between the conceptual/semantic slots of a concept and the syntactic slots of a linguistic expression? At least with respect to the list of expressions I have been dealing with so far, I think that this distinction makes sense. My aim in this section is that of showing that it makes sense to separate the *syntactic* aspects and the *semantic/conceptual* aspects of expressions like "rain", "ready", "green", etc.: the semantic/conceptual structure of these expressions can be shown as richer than their syntactic structure.

Evidence that this is the case is given by cases in which the predicate (whether "rain", "ready" or "green") is correctly used with a monadic syntactic structure, whereas clearly the concept (meaning) associated with the predicate has a more complex (though implicit) structure. In **Chapter 3** I

introduced the notion of "zero-reading", i.e. of a reading in which the predicate is correctly used with a monadic syntactic structure. The sentence

(10) Whenever the temperature and humidity in the air reach a certain level, it rains.

is such that the predicate "to rain" that occurs in it can get a zero-reading, i.e. a reading where "to rain" expresses a monadic property of times. Analogously, the sentences

(11) Whenever she is prepared for an activity, Jill is ready
(12) Whenever something or some salient part of it looks like this, it is green

contain predicates for which a zero-reading is appropriate, i.e. in which "being ready" and "being green" can be interpreted as monadic properties of objects. That these predicates can be correctly used with a monadic syntactic structure doesn't imply that the concepts (meanings) we associate with them do not have a more complicated structure. Clearly the concept TO RAIN doesn't apply just to times; it concerns also places, substances (e.g. water, sulphuric acid), manners in which the phenomenon may manifest itself. The word "to rain" in English, as well, seems to have a meaning whose semantic complexity is not exhausted by the monadic structure of the predicate *rain(t)*. All this speaks in favour of a greater complexity of the concept TO RAIN (and of the meaning of "to rain"), as opposed to the syntactic structure of the predicate "to rain". These considerations obviously apply to "ready" and "green" as well.

I propose to rephrase the distinction between argument-places and conceptual (semantic) constraints as a distinction between *argument-places* and *adjuncts*: what I call conceptual (semantic) constraints find a representation as adjuncts, and not as argument-places.

Roughly, an *argument-place* is a characteristic feature of the logical structure of an expression. Of expressions which have argument-places, Frege said that they are "unsaturated entities" (see Frege 1892). Thus, if an expression contains an argument-place, it will not give a determined contribution to the proposition expressed by a sentence in which it occurs until that argument has been assigned a value, or "saturated". This means that providing the value of the argument reduces the degree of unsaturatedness of the expression, i.e. it reduces its "adicity". An *adjunct*, on the contrary, does not cause an unsaturated expression to become saturated, i.e. it doesn't decrease its adicity. It merely acts as a modifier of the expression. As Heim and Kratzer write:

> Arguments reduce the adicity of the [verb] they combine with; modifiers leave it unchanged. [...] Restrictive modifiers are characterised by the fact that they leave the semantic type, as well as the adicity, of the modifiee completely unchanged. (Heim and Kratzer 1998: 64).

Adverbs are a typical case of modifiers. Take, for instance, the adverb "continuously": when added to a verb like "stumble", it has the effect of making the truth-conditions of a sentence in which the verb occurs more precise or fine-grained: there is a difference between the truth-conditions of "Sarah stumbles" and the truth-conditions of "Sarah stumbles continuously". In a way analogous to adverbs, adjuncts are devices of modification and not of saturation. For example, in the sentence "Sarah sneezes in the living room" the phrase "in the living room" is an adjunct in the sense that it modifies the predicate by intersecting the extension of "x sneezes" with the extension of "x is in the living room"; it could be thought of as specifying an aspect of the event of Sarah's sneezing, which pertains to the location of the event. Note that this doesn't either entail or require that the expression "to sneeze" has a hidden argument in deep syntax for the location of the sneezing to be saturated.

Having introduced the distinction between arguments and adjuncts, the point I wish to make is the following: on the one hand, argument-places determine mandatory constraints on sentential well-

formedness. On the other hand, adjuncts, which are a way of representing conceptual (semantic) constraints, do not determine mandatory constraints on sentential well-formedness.

To illustrate the difference, consider a predicate like "give", which expresses a triadic relation *give(x, y, z)* between a giver, a receiver and a given item. Suppose that a speaker utters the following in a context in which no particular given item is presupposed.

(13) John gives Mary;

Sentence (13) is not well-formed and it doesn't express a fully intelligible content. The reason is that the speaker is, as it were, disregarding the syntactic structure of the predicate; namely, she is failing to fill in the predicate's argument place corresponding to the given item. (Some may point out that (13) doesn't sound so bad. However, I suspect that that is because there are uses of "give" in which the argument-place for the given item is controlled by a (silent) existential quantifier, the resulting content being *<that John gives (something) to Mary>*. In these cases, the predicate's structure is upheld, which explains why there's no intuition of bad formation.) Contrast this with the following situation, in which the speaker utters:

(14) Mary sings

Even though the minimal content *<that Mary sings>* is not truth-evaluable—after all, if Mary sings, she necessarily sings something—this by no means makes (14) non well-formed and unintelligible. (14) is a perfectly well-formed sentence. In this case, the speaker failed to articulate an adjunct for the sentence, i.e. an element that could have modified the content of (14) by making it more precise, or specific (in this case, also fully propositional); she didn't however violate any *grammatical* rule about the logical/syntactic structure of the word "sing".

Thus, to understand the notion of conceptual constraints I have been concerned with so far, the notion of adjunct is of help insofar as adjuncts specify aspects of the conceptual (semantic) structure of predicates without thereby being ways of "saturating" any argument-place articulated in the syntactic structure of the predicates themselves.

One thing that I may not have emphasised enough is that *both* argument-places and adjuncts express conceptual (semantic) constraints. After all, argument-places could be thought of as articulating the *most powerful or fundamental* conceptual (semantic) constraints, like for example the constraint that the "giving" relation necessarily obtains between a giver, a receiver and a given item! With this qualification in place, it seems safe to conclude from the previous considerations that the conceptual constraints articulated by argument-places determine constraints on well-formedness (and intelligibility), while those conceptual constraints articulated by adjuncts do not determine constraints on well-formedness (and intelligibility). This I take as a viable way of making sense of the difference between an argument-place, which is supposedly articulated at the syntactic level, and a simple "conceptual constraint", which needs not be syntactically articulated but is nevertheless semantically/conceptually operative.

Going back to the contrast between a syntactic approach and a conceptual constraints approach to comprehension, one may claim that the argument-place vs. adjunct distinction is of help in setting the two proposals apart. According to the proponent of the syntactic approach, many of the aspects that the conceptual constraints view regards as captured by adjuncts are really captured by genuine, syntactically realised argument-places. This implies that the logical form of many predicates is treated as *unexpectedly complex*: for example, "to rain" is treated as if it was a dyadic relation between times and locations; "being ready" is treated as if it was a dyadic relation between individuals and purposes, and so on. This view faces a serious challenge from over-generation, for it seems that there is no principled way of ruling out that further hidden argument-places may be "discovered" (see Cappelen and Lepore 2002, 2005, Hall 2008, whose arguments I considered in

Chapter 3, **section 5**). By contrast, according to the conceptual constraints view, we can retain a conception of these predicates as monadic at the syntactic level, while allowing a greater complexity at the conceptual/semantic level, and this is because we have the notion of adjunct (or modifier) at our disposal.

To sum up, in this section I have highlighted the difference between a syntactic approach, which explains utterance comprehension in terms of the satisfaction of hidden argument-places in deep syntax, and the conceptual (semantic) constraints account, which explains utterance comprehension by making no commitments as to the syntactic implementation of such constraints. The proponent of the conceptual constraints view does take a commitment with respect to the idea of there being conceptual structures determining a *standard cognitive (discoursive) format*. Interpreters of utterances of semantically under-determined sentences such as "It's raining" may be portrayed as carrying out their interpreting inference by constantly referring to a standard cognitive (discoursive) format with which the content of the speaker's utterance is expected to comply. The strive to a standard "format" for utterance content is not equivalent to compliance with syntactic constraints, such as argument-place saturation. There is a difference between satisfying the syntactic constraints of a predicate and satisfying the conceptual constraints of that predicate. The difference may be put as follows. At a syntactic level, mandatory constraints on well-formedness are determined; conversely, at a conceptual (semantic) level no such mandatory constraints are determined.

6. Conceptual Constraints and the Semantics/ Pragmatics Distinction

So far I have been describing the details of how conceptual constraints can help hearers in the reconstruction of the content of utterances of certain semantically under-determined sentences, such as "It's raining", "Jill is ready", "The leaves are green", and so on. I have been concerned with setting the proposal apart from already familiar options, such as inferentialist, non-inferentialist and syntacticist options, but also from other explanations in terms of conventional or generalised implicature. After having set the boundaries of the view, a residual question may be raised. Do the processes that I have characterised belong to the realm of *semantics* or to the realm of *pragmatics*? The question appears altogether legitimate, given the relevance that has been awarded to the distinction by the recent debate in the field (see e.g. Turner 1999, Szabó 2005). The issue, moreover, is not merely terminological, i.e. it is not entirely about deciding which label to attach to which phenomena; as I see it, it is about deciding how much of the non-natural product that we call "meaning" (see Grice 1957) depends *solely* on the intentions of the speaker and how much depends on the intentions of the speakers *plus* some sort of conventions, or rule-governed practices. The portion of meaning which depends solely on the speaker's intentions falls under the heading of "pragmatics", while the portion which depends on speakers' intentions plus conventions falls under the heading of "semantics".

Having set out these considerations, let us turn to the issue whether the processes I have been describing so far belong to the pragmatic rather than to the semantic realm. There is more than one way of approaching the distinction between semantics and pragmatics. Philosophers like Grice, Austin and Searle initiated what we may call the *Received Conception* about the semantics-pragmatics distinction, which is characterised by a *sharp separation* between the two domains. Among them, only Grice's work will be relevant for our purposes[50]. Grice ((1989), see especially the essays "Logic and Conversation", "Further Notes on Logic and Conversation" and "Utterer's Meaning, Sentence Meaning, and Word Meaning") is interested in the semantics/pragmatics distinction insofar as it can illuminate a distinction between *conventional* and *intentional meaning*. The Gricean view has it that semantics concerns the truth-conditions of an utterance, which in turn

---

[50]   Grice's interests and Austin's and Searle's are not exactly the same. Austin (1962) and Searle (1969, 1975) are motivated by the purpose of distinguishing between different *kinds of acts*. Their distinction is between "merely saying" that *p*, i.e. uttering a sentence *s* whose content *p* depends on semantic conventions (*locutionary* act), and "doing something while saying" that *p*, namely asserting, commanding, promising, etc. (*illocutionary* act).

are thought of as determined by the *conventional* content of the utterance's components and by their syntactic arrangement[51]. Pragmatics, on the other hand, deals with the non-truth-conditional aspects of an utterance, i.e. with what an utterance does not say, but rather "conveys" or "implicates": these aspects are determined by the communicative intentions of the utterer. Followers of Grice and supporters of his version of the *Received Conception* are, among others, Bach and Harnish (1979), Cappelen and Lepore (2005), Soames (2002), Borg (2004).

The Gricean Received Conception has been challenged by what may be called the *Pragmatist Conception* (see Levinson 2000, Recanati 2004, Carston 2004, Travis 2008, Korta and Perry 2007), i.e. the view that there is no sharp separation between semantic and pragmatic aspects. According to the subscribers of the Pragmatist Conception, what an utterance of a sentence "says" in context cannot be equated with its semantic content, i.e. the result of the composition of the semantic contents of the utterance's constituents, even conceding the role of context in disambiguation and in fixing the reference of indexicals and demonstratives. Mechanisms different from formal semantic, compositional rules, and even from those rules related to indexicality, affect the content of an utterance in context. These are typically called mechanisms of "free enrichment": they affect the content of an utterance "top-down" rather than "bottom up", that is, they respond to pragmatic demands rather than to syntactico-semantic demands. Processes of "free enrichment" determine a mutual contamination between semantic and pragmatic aspects, between "what is said" and "what is implied" (bringing about what Levinson calls "Grice's circle"). This brings to a redefinition of the boundary between semantics and pragmatics, according to which the truth-conditions of an utterance are determined by *both* semantics and pragmatics, while the non-truth-conditional aspects of an utterance are still dealt with entirely in pragmatic terms.

My aim in this section is to argue that conceptual constraints definitely belong to semantics rather than pragmatics, however they are aspects of the semantics of terms that are not strictly speaking articulated or constrained at the syntactic level. Ultimately then, the conceptual constraints view endorses the *Received Conception* in recognising a sharp separation between semantics and pragmatics, with the *proviso* that certain semantic mechanisms are not constrained by syntax.

The effects produced by conceptual constraints are arguably semantic rather than pragmatic. There is more than one reason to think so. First of all, conceptual constraints can legitimately be considered semantic constraints for any speaker who is competent on the meaning of an expression *e* of *L*. For instance, the conceptual constraints on RAIN are also semantic constraints on the predicate "to rain" in English, for any speaker who masters this predicate. The constraints are "semantic" insofar as they sanction correct use of the words at interest: for instance, the conceptual constraint that pertains to rain locations can legitimately be regarded as a "semantic" constraint, to the extent that instructions can be derived from it that regulate the correct use of the verb "to rain". Secondly, conceptual constraints can legitimately be regarded as semantic because they operate on what is said by an utterance, rather than on what is implied. Again, the rain-location constraint arguably generates effects on the truth-conditions of utterances of "It's raining", rather than determining a further level of content, such as an implicature. That conceptual constraints operate on what is said and hence on the semantics of an utterance, rather than on what is implied by it, appears from the fact that their effects on the utterance's content cannot be consistently cancelled (see **section 3.1** in this chapter). For instance, assuming that conceptual constraints allow a hearer to go from the content *<it's raining>*, to the content *<it's raining somewhere>*, it would be contradictory for the speaker to cancel the effect of such constraints by asserting "It's raining, but nowhere". This wouldn't be an implicature cancellation: it would be a sheer contradiction. The

---

[51]   In a Gricean framework, the references in context of indexicals ("I", "here") and demonstratives ("he", "this") are to be considered fully part of the semantic content of an utterance. A "minimal" amount of context-sensitivity is therefore tolerated, as long as it can be demonstrated that an expression is assimilable to either pure indexicals or demonstratives. Bach (2002) sets forth for a diverging view, which expunges the content of demonstratives, but not that of "pure" indexicals from semantic content.

inconsistency that results from an attempt to cancel the effects of conceptual constraints I therefore take as a strong indication that they trigger semantic, rather than pragmatic mechanisms.

Conceptual constraints may be thought of as opening "semantic slots", which in appropriate occasions may be filled in with contextual information. For instance, the rain-location constraint for "to rain" may be regarded as opening a semantic slot to be filled with a location (Paris, London, the San Francisco Bay, etc.). Saturation of these slots needs retrieval of contextual information, perhaps by inquiring into the speaker's intentions, beliefs and so on. Only after the intended location has been identified, the location-slot may be filled in. Notice, however, that this doesn't count as a pragmatic process any more than other processes of reference-assignment to linguistic expressions (such as "this", "that") on the basis of contextual information count as pragmatic. Since the location-dedicated slot is opened up for semantic/conceptual reasons (which relate to a "standard cognitive/discoursive format"), the saturation of the slot is analogous to the saturation of any other free argument-place through contextual variable assignment.

Conceptual constraints therefore determine semantic, rather than pragmatic effects. At least with respect to the semantic profile of sentences such as "It's raining", "Jill is ready", "The leaves are green", the *Received View* may be preserved, that is to say, a sharp separation may be maintained between the semantic aspects of the content-in-context of an utterance of "It's raining" *vis-à-vis* its pragmatic aspects. However, notice that the distinction is to be supplemented with the following *proviso*: among semantic effects, some are going to be syntactically *un*constrained. For instance, the effect that determines enrichment of the content *<that it's raining>* as *<it's raining somewhere>* doesn't follow any syntactic restriction, i.e. it is not dictated by the syntactic structure of the predicate "to rain", but only by conceptual (semantic) constraints on correct use of the concept/expression. The same holds for predicates such as "being ready" or "being green": the fact that semantic processes allow one to go from *<Jill is ready>* to *<Jill is ready for something>* doesn't imply that these processes are syntactically regulated; analogously, from the fact that the step from *<the leaves are green>* to *<the leaves are green in some respect>* is semantic in nature, it doesn't follow that it is syntactically controlled.

In conclusion, the Received View about the semantics-pragmatics distinction is maintained with respect to sentences such as "It's raining", "Jill is ready", "The leaves are green", ec., with the qualification that there are semantic processes (i.e. processes on what is said based on conceptual/semantic aspects) which however are syntactically unconstrained. The following chart illustrates the idea:

Formal Semantics	Conceptual/Semantic Constraints	Pragmatic processes
What is said Syntactically constrained	What is said Syntactically *un*constrained	What is implied Syntactically *un*constrained

Conclusion

This chapter has been devoted to a clear statement and some refinements of my proposal as to the role of conceptual constraints in the comprehension of utterances of semantically under-determined sentences. Conceptual constraints are ways of structuring thought, language use and behaviour in general, and they provide instructions as to what a concept-user should do in order to correctly apply a concept. For example, that every event of raining is spatially located; that every instance of readiness is readiness for something; that every instance of greenness is greenness in some way or respect are all constraints on the concepts of RAIN, of READINESS and GREENNESS. My claim was that use of these conceptual constraints helps hearers to determine a content for utterances of semantically under-determined sentences such as "It's raining", "Jill is ready", "The leaves are green", etc. The chapter has been structured in the following way:

- I have delimited the scope of my proposal, stating what I am and what I am not concerned with. I am *not* concerned with an account of non-sentential assertions, or with an account of evaluable, but pragmatically inaccurate assertions. What I *am* concerned with is utterances of sentences that are semantically under-determined in that they do not express an evaluable content for their utterances, even though they do not immediately strike as examples of ambiguity, indexicality, vagueness, ellipsis;
- I have presented some of the main, current approaches to utterance comprehension. I reviewed three accounts: (i) an *inferential account*, according to which utterance comprehension (especially of semantically defective expressions) is a prominently inferential task. The inferential position comes in two varieties: a Gricean variety, countenancing explicit inferences; and a Relevance-theoretic variety, countenancing implicit or sub-personal inferences performed by dedicated mental modules; (ii) a *non-inferential account*, according to which the comprehension of the content of certain utterances is as direct as perception or mediated by associative processes, and thus non-inferential; (iii) a *syntactic account*, according to which the pragmatic effects that are invoked in explaining comprehension of a semantically under-determined sentence are all traceable to hidden positions articulated in logical form;
- Against non-inferentialist views, I argued that the contents of semantically under-determined sentences, though semantically defective, are definitely conceptually structured, and hence can participate in inferential relations;
- Against the Gricean inferentialist view, I defended the implicit character of interpretive inferences. Following Sperber and Wilson, I endorsed the view that interpretive inferences are performed sub-personally, in a fast a painless way. Against Sperber and Wilson, though, I have argued that the interpretation of such utterances as "Jill is ready", "It's raining", etc. *does not issue in a meta-representational conclusion*. Rather, the process may consist of an inference whose premises and conclusion are entirely *meta-representation-free*. The role of meta-representation has to be reconsidered: meta-representational capacities should be seen as employed in *justifying* some of the premises of the interpretative inferences, and not as *providing* any of these premisses;
- I have claimed that the syntactic approach, which takes commitments as to the syntactic underpinnings of a large part of processes of utterance understanding, doesn't argue for its syntactic commitments in a sufficiently cogent way. In contrast, the conceptual constraints view doesn't take any syntactic commitments as to the presence of covert components. Utterance comprehension is guided by a strive to a *standard cognitive (discoursive) formatting*, rather than by genuine syntactic constraints. The difference between syntactic constraints and conceptual constraints may be cashed out as tha difference between *argument places* and *adjuncts*: while argument places pose constraints on well-formedness, conceptual constraints do not pose any such constraint;
- I have argued that the application of conceptual constraints gives rise to semantic processes, which are however not syntactically constrained. The conceptual constraints view therefore sticks to the Received Conception of the semantics-pragmatics distinction, which envisages a sharp separation between the two domains, however it endorses the thesis that *there may be semantic, not syntactically constrained processes.*

Conclusion

Let me conclude this work by summarising the two main claims that I have defended in this thesis, as well as by sketching some possible lines of future developments. I started out this thesis by acknowledging—as many theorists do—that linguistic meaning generally under-specifies the truth-conditions of utterances: this was the *Under-Specification Claim*. Within the variegated landscape of semantically under-specific expressions, I have isolated some expressions which I have called for convenience *semantically under-determined*. These are under-specific expressions whose semantic defectiveness does not immediately allow a reduction to familiar forms of under-specification, such as ambiguity, vagueness, indexicality, ellipsis. Some examples are: "It's raining", "Jill is ready", "Lisa is tall", "The leaves are green".

1. The Under-Articulation Claim

The first problem with which I have been engaged is that of distinguishing semantic under-determinacy from other phenomena of semantic defectiveness—or under-specification. I have called this problem *The Demarcation Problem*, which could be summarised by the question "Is there a feature which sets semantic under-determinacy apart from ambiguity, vagueness, indexicality, ellipsis, etc.?". As we have seen, many theorists agree that semantic under-determinacy is a *sui generis* phenomenon which should not be "domesticated", or reduced to any already known kind of under-specificity. At a first glance, semantic under-determinacy seems to have nothing to do with ambiguity, with vagueness, or with indexicality (see Travis 1997, Bezuidenhout 2002). This I take as being a negative characterisation of semantic under-determinacy—corresponding to a claim as to what it is *not*. My aim is to say something positive about semantic under-determinacy, and I believe that providing an answer to the Demarcation Problem is a promising strategy in this direction. In order to give a clear response to the Demarcation Problem, it seems appropriate to explore the question of where semantic under-determinacy comes from—of what its *source* or *cause* is. It is by answering this question that, I argue, it shall be possible to demarcate semantic under-determinacy with respect to other familiar phenomena of semantic under-specification. **Chapter 1** is devoted to the problem of the source of semantic under-determinacy, as a way of dealing with the Demarcation Problem.

On the one hand, it seems incorrect to say that semantic under-determinacy comes from the fact that *meaning itself* is indeterminate: when a speaker uses a semantically under-determined expression, like "Jill is ready" there is a clear sense in which anyone can understand what she is saying on the account that the speaker is using that expression for what it means, i.e. for what its public, conventional meaning is. So, there is a clear sense in which the meaning of semantically under-determined expressions is fixed and determinate. One may then suppose that what is indeterminate is the *content-in-context* of semantically under-determined expressions; this is probably a better way of putting the matter. However, if we are ready to accept that semantic under-determinacy depends on the indeterminacy of the content-in-context of the expression's utterances, then semantic under-determinacy becomes indistinguishable from indexicality (context-sensitivity in general). After all, the hallmark of indexicality is precisely that, although the meaning of an indexical expression is determinate, its content-in-context is indeterminate.

Indexicality and semantic under-determinacy are different. The best way to make this difference apparent is to look at the differences between indexicality-resolution and the resolution of semantic under-determinacy. Capitalising on considerations that were already advanced more or less explicitly by Searle and Travis, my argument for the difference between the two phenomena is the

following:

[The Reiteration Effect Argument] Indexicals are designed in such a way that, once the indexicality of an expression *e* is resolved in a certain context, the semantic under-specificity of *e* is resolved in all circumstances. By contrast, once the semantic under-determinacy of an expression *e* is resolved in one context, its under-specificity may *not* be resolved in all circumstances. It may be "reiterated" in other circumstances. If this is so, then semantic under-determinacy and indexicality cannot be the same phenomenon.

The fact that semantic under-determinacy can be distinguished from indexicality on the basis of its giving rise to the "Reiteration Effect" is a step towards a characterisation of the phenomenon, but it's not a conclusive result yet. For the "Reiteration Effect" is just an "effect", i.e. it is a consequence of semantic under-determination. What has to be found out is what *gives rise* to this "effect".
My proposal concerning the source of the Reiteration Effect is to "move one level down" from the stage of semantic phenomena to that of *linguistic* phenomena. As a preliminary observation, I remark that, for reasons of economy, brevity and effectiveness, speakers often leave their thoughts "unspoken" to a certain degree. For example, if all the conversationalists know that the subject of their discussion is the weather conditions in Paris, one might just utter "It's raining" rather than "It's raining in Paris"; if all know that the topic is Jill's preparing her Spanish exam, one might just utter "Jill is ready" rather than "Jill is ready for the exam". What I have described informally as "leaving things unspoken" may be recast in a more precise way in terms of producing *under-articulated sentences*, where the under-articulation of a sentence is thus defined:

[Under-Articulation] A sentence *s* is under-articulated iff, when an utterance *u* of *s* is performed, another sentence *s'* could have been uttered instead of *s*, which articulates more linguistic material than *s* and has the same (intended) truth-conditions as *u*.

In other words, a sentence *s* is under-articulated iff the (intended) truth-conditions expressed in a context *c* by its utterance *u* could have been expressed by a more articulated sentence *s'*. Having given this definition of under-articulation, I move on to considering the claim that semantic under-determinacy arises precisely from under-articulation.

[The Under-Articulation Claim] The meaning of a sentence *s* under-determines the (intended) content and truth conditions of its utterance *u* because and only because *s* is under-articulated.

The main advantage of the view that stems from the Under-Articulation Claim is that it allows to give a clear answer to the Demarcation Problem, for it seems that cases of semantic under-determinacy all enjoy the following feature: they are *explainable solely in terms of linguistic under-articulation*. For example, the under-determinacy of the sentence "It's raining" as uttered by Eva while meaning *<that it's raining in Paris>* might be explained by solely appealing to the fact that the sentence itself "It's raining" is under-articulated with respect to a sentence which linguistically articulates the content which Eva intends to convey—perhaps the sentence "It's raining in Paris". In general, cases of semantic under-determinacy may be explained solely in terms of under-articulation. This contrasts with the way ambiguity, vagueness, indexicality may be explained: in neither of these cases it seems that the phenomenon is exhaustively explained in terms of under-articulation only. Ambiguous expressions aren't such only because of under-articulation, but because words can get more than one meaning; vague expressions aren't such only because of shortage of linguistic material—at the source of vagueness is the fact that the application of certain predicates is unclear in some cases; indexical expressions aren't such only because we leave things unspoken—at the root of indexicality is mechanisms of context-exploiting in order to determine the

174

referent of a restricted group of terms. Note that under-articulation relevantly differs from ellipsis, in that under-articulated sentences can be felicitously uttered even if no sentence has been articulated previously, contrary to what happens to elliptical expressions.

Having provided an answer to the Demarcation Problem, the next problem to be confronted is a perhaps more pressing one: *how is communication possible, despite the fact that semantically under-determined sentences are used by speakers?* In the thesis, I tackle this problem by proposing a view in which successful communication via under-determined sentences is possible in virtue of speakers and hearers largely relying on what I call *conceptual constraints*.

2. The Conceptual Constraints View

Speakers utter myriads semantically under-determined sentences every day. Hearers manage to somehow understand what they say. How is that possible? As one may notice, the big problem of communication with semantically under-determined sentences is primarily a *comprehension* problem. Under-articulation is a hurdle to understanding what a speaker means with a certain sentence. That is tantamount to say that semantic under-determinacy is a problem for hearers and not for speakers, who presumably know what they mean even when they leave some words unspoken. The problem could be spelled out as:

[The Comprehension Problem] What is the general mechanism by which the listener of an under-determined sentence *s* acquires comprehension of what the speaker means by uttering *s*?

In the thesis, I scrutinise various proposals inspired by the most popular theories of the relation between language and context. The answer inspired by Radical Contextualism has it that *pragmatic processes* are necessary in order to achieve comprehension of utterances of semantically under-determined sentences. The answer inspired by Indexical Contextualism is to the effect that the comprehension of semantically under-determined sentences necessarily depends on the resolution of covert forms of *indexicality* (e.g. saturation of hidden slots in logical form). Finally, the answer that a Truth Relativist could give has it that comprehension of semantically under-determined sentences necessarily requires recognising a form of *truth relativity*. In **Chapters 2**, **3** and **4** I provide reasons for being dissatisfied with each of these answers. In general, I am persuaded that the comprehension of semantically under-determined sentences is exhaustively captured neither by an entirely pragmatic account (such as the Radical Contextualist Account) nor by a purely syntactic-formal account (such as the Indexical Contextualist and the Truth Relativist Account).

My proposal is to account for the comprehension of semantically under-determined utterances by appealing to the idea of *conceptual constraints*. Conceptual constraints are primarily *cognitive tools*, they are ways of organising thought on specific matters. Employing a pictorial representation, we may regard conceptual constraints as *schemata*, where a schema is a structured configuration in which various elements are displayed as standing in various relations to each other. For example, we seem to hold a *schema* for thinking about events of raining: an event of raining is a physical phenomenon (the falling of drops of water from the sky), which can stand in relation with times and with locations. Another *schema* is that for thinking about readiness: being ready is a state which involves a subject (of which readiness can be predicated) as well as some purpose or activity. And so on. An important qualification is the following: Conceptual constraints *coincide* with semantic constraints for all competent language users: for instance, one who masters the meaning of "to rain" in English is expected to be competent also on the concept of RAIN. My inclination is however that of privileging talk in terms of conceptual constraints insofar as I deem the conceptual domain more stable and reliable than the semantic domain, which depends on language mastery. I therefore endorse the following claim:

[The Conceptual Constraints Claim] Conceptual constraints play a role in the comprehension of

semantically under-determined sentences—namely, they are employed by the hearer as a guide for reconstructing their content and truth-conditions.

So for example, suppose that a speaker utters the sentence "Jill is ready", thereby meaning that Jill is ready to ski. The hearer, who we may suppose is aware that the conversation topic is Jill's being about to ski, may engage in the following (implicit) reasoning. After having decoded the sentence's meaning *<that Jill is ready>*, she may apply the conceptual (semantic) constraint to the effect that being ready involves a purpose or activity; knowing from context that the salient purpose is skiing, she may derive from this the content *<that Jill is ready to ski>*.

What is interesting about this reasoning is that it is driven by conceptual (and in this case, also semantic) constraints; that is, it is competence about the conceptual (semantic) restrictions that operate on uses of "being ready" that allows the hearer to go through this simple reasoning. Some interesting questions arise about conceptual constraints and the processes of comprehension that they may be taken to govern. Do conceptual constraints determine semantic or pragmatic processes? What kind of abilities are involved in these processes?

I have argued that, on the one hand, processes driven by conceptual constraints definitely belong to the realm of *semantics*, because they impose restrictions on aspects of the truth-conditions of the utterance performed, and not on any other content that may be merely suggested or implied by the utterance—like pragmatic processes. Nevertheless, conceptually constrained processes are not governed by the syntax of the expression involved. In other words, conceptually constrained processes determine semantic effects (which affect "what is said" by an utterance) which are not subject to syntactic restrictions. The following chart illustrates the relation between the semantics of utterances, syntax, conceptual constraints and pragmatics.

Syntactic constraints	What is said by an utterance (Semantics)
Conceptual (semantic) constraints	
Pragmatic principles, maxims	What is implied by an utterance (Pragmatics)

Concerning the abilities involved in the performance of conceptually constrained processes, I have argued that these processes can be framed as entirely *meta-representation free*, i.e. as consisting of inferences whose premises and conclusion involve no attitude-attribution to the speaker such as "X means that P", "X believes that P", "X intends that P". By so doing, I have contrasted my proposal with two very influential models of utterance comprehension, which both conceive of this process as being intrinsically meta-representational: the Gricean model, according to which utterance comprehension is meta-representational through and through; and the Relevance-Theoretic model, according to which utterance comprehension gives rise to at least meta-representational conclusions.

3. Future developments
So far, we have invoked an account of utterance comprehension that relies on conceptual constraints for a class of expressions whose content could be characterised as, among other things, "conceptually truncated" (see Bach 1994a). These are sentences such as "It's raining", "Jill is ready" which could be described as missing a complement, such as "... in Paris", "... for skiing". It's easy to draw a parallel between the possible complements for these sentences ("... in Paris", "... for skiing") and the corresponding elements in the proposition that are dictated by conceptual constraints *plus* contextual information (*<Paris>*, *<skiing>*). In other words, it's very straightforward to invoke conceptual constraints where the semantic defectiveness of the expression seems to arise, among other things, from some sort of *conceptual* incompleteness of its content.

This being said, it seems that conceptual constraints and the related cognitive *schemata* may also play a role in facilitating the comprehension of expressions whose content is not conceptually incomplete. Here *schemata* may be more complex configurations representing types of situations that are highly relevant for a certain social or cultural group. For example, consider the sentence "John was arrested by a policeman yesterday; he had just stolen a wallet". Recanati (2004) argues that the stealing-and-being-arrested *schema* affects reference assignment to the pronoun "he" in the second clause, therefore conditioning the truth-conditional interpretation in the sentence. As he observes:

> John is selected as the referent for 'he' in the second clause because (i) the referent of 'he' is said to have stolen and John is known to have been arrested, and (ii) there is a frame or schema in which the two roles (stealing, and being arrested) are linked. This schema is jointly activated by the predicates 'was arrested' and 'had stolen'. An interpretation in which the same person steals and is arrested (and in which he is arrested because he has stolen [...]) satisfies the schema, and is more likely to be selected than one which violates it. (Recanati 2004).

Similarly, one could suppose that *schemata* and the related conceptual constraints play a role in the interpretation of those expressions that give rise to "*ad hoc* concepts", i.e. temporary, context-dependent and goal-directed conceptual constructions (see Barsalou 1983). Consider for example the sentence "John took the key and opened the door". It seems that a particularly favourable interpretation of this sentence is true just in case John took the key and opened the door with the key. Here one may suppose that the first conjunct in the sentence activates a use-the-key-to-open-doors *schema*, which affects the interpretation of the predicate "open" in the second conjunct. The occurrence of "open" is taken to denote the "ad hoc" concept OPEN-WITH-THE-KEY rather the more generic and not *ad hoc* concept OPEN. In a similar vein, one could suppose that the interpretation of a compound expression like "dog magazine" be guided by some complex *schema*-like representation: in working out the *ad hoc* concept denoted by this compound locution, one may invoke a sub-schema relative to magazines which envisages magazines-devoted-to-a-specific-topic, and therefore conclude that the expression denotes the concept MAGAZINES-ABOUT-DOGS (see Murphy 2002).

In light of these considerations, it might be interesting to explore the possibility of there being a *continuity* between the conceptual constraints that drive the comprehension of sentences with conceptually defective contents such as "It's raining" and "Jill is ready" and the constraints involved in the interpretation of the non (clearly) conceptually-defective contents just surveyed. The first kind of conceptual constraints seems to be more deeply entrenched in our way of thinking about the world: after all, it's not possible to conceive of rain-falling events as non spatially located, or of states of readiness as not related to purposes (though, as I argued in **Chapter 3**, it is certainly possible to felicitously use expressions like "It rains" while referring to no specific place). The second kind of conceptual constraints seem to arise from the way we conceptualise certain specific matters, such as stealing and its punishment, key-usages and thematic publications. These *schemata* seem to be less entrenched in our conception of the world, in that they strongly depend on the way we organise our habits, our cultural and social practices: we use keys to open doors, we (usually) arrest those who steal and publish thematic magazines. Having recognised these important differences, the possibility remains that, with respect to utterance comprehension, these *schemata* function much in the same way, i.e. they are resources that hearers employ in reconstructing the content and truth-conditions of certain utterances.

A further issue that I deem particularly interesting concerns the role of meta-representation in those processes that I regard as "conceptually constrained". As I stressed in **Chapter 5** I believe that, once we recognise the role of conceptual constraints in the comprehension of utterances of "Jill is ready",

"The leaves are green", "It's raining" etc. the role of meta-representation has to be substantially revised. Hearers may use their capacity to meta-represent in order to fix some of the utterance's *background aspects*, such as what the speaker is talking about, what she is making reference to, etc. However, they do not need to carry out the whole comprehension process in meta-representational terms. The most remarkable consequence of this idea is that the upshot of processes of utterance-comprehension that involve semantically under-determined sentences needs not be a meta-representation, as for example an attribution of meaning ("X means that P"), of intention ("X intends that P"), or belief ("X believes that P"). In a nutshell, our interpretive conclusions need not be meta-representations: they might as well be object-level representations.

One interesting route for the theorist willing to pursue this line of investigation is that of devising empirical tests aimed to prove that the result of the comprehension processes for this group of sentences is at the object-level rather than at a meta-level. The basic idea is that of asking subjects to report what they understand, after having read or heard a brief story which culminates with an utterance of a semantically under-determined sentence such as "It's raining". For instance, subjects may be asked to read/hear the following story:

> "Suppose you have a friend who lives in Paris. You often call each other on Skype, and you often begin the conversation by exchanging information about the weather in each of your locations of residence. Suppose you are on one of these Skype calls and you ask your friend: "So, what's the weather in Paris?"; she answers "Not quite good. It's raining." What information can you say to have gathered from her answer?".

If the conceptual constraints account is on the right track, the prediction is that *subjects will tend to give non-meta-representational answers*, like "That it's raining", "That it's raining in Paris", "That the weather is not quite good, since it's raining where she lives", rather than meta-representational answers such as "That she means that it's raining in Paris", "That she believes that it's raining in Paris", and so on. If it were so, this would provide evidence that meta-representational information, rather than being the result of the interpretation process of an utterance of "It's raining", is at least an intermediate step, one at which some important background information as to the intentions of the speaker is settled. This meta-representational, background information supports an interpretive result which is, however, *non*-meta-representational. I believe there are good prospects for experimental results that confirm the prediction, though this will inevitably have to be established by future research.

Bibliography

Austin, J. L. 1950. Truth. *Proceedings of the Aristotelian Society,* Suppl. Vol., 24: 111–128.

Austin, J. L. 1962. *How to Do Things with Words*, Oxford: Clarendon.

Bach, K. 1994a. Conversational Impliciture, *Mind and Language*, 9: 124-62.

Bach, K. 1994b. Semantic Slack. In Tsohatzidis, S. L. (ed.) *Foundations of Speech Act Theory*. Routledge: 267-91.

Bach, K. 2000. Quantification, Qualification and Context a Reply to Stanley and Szabó. *Mind and Language* 15 (2&3): 262–283.

Bach, K. 2002. Seemingly Semantic Intuitions. In J. Campbell, M. O'Rourke, and D. Shier, eds., *Meaning and Truth: Investigations in Philosophical Semantics*. New York: Seven Bridges Press: 284-292.

Bach, K. 2012. Context-Dependence (Such as it is). in Garcia-Carpintero, M. and Kölbel, M. (eds). *The Continuum Companion to the Philosophy of Language*.

Bach, K., and Harnish, R. 1979. *Linguistic Communication and Speech Acts*, Cambridge, MA: MIT Press.

Baker, C. 2012. Indexical contextualism and the challenges from disagreement. *Philosophical Studies*, 157:107–123.

Bar-Hillel, Y. 1954. Indexical Expressions, *Mind* 63, pp. 359-379.

Barba, J. 2007. Formal semantics in the age of pragmatics. *Linguistic and Philosophy* 30: 637–668.

Barsalou, H. 1983. Ad Hoc Categories. *Memory & Cognition*, 11: 211-227.

Barwise, J. & Perry, J., 1983, *Situations and Attitudes*, Cambridge, MA: The MIT Press.

Bealer, G. 1998. Intuitions and the Autonomy of Philosophy. in DePaul M.R. and Ramsey W. (eds.) *Rethinking Intuition*, Lanham, MD: Rowman & Littlefield: 201-39.

Bezuidenhout, A. 2002. Truth-conditional Pragmatics, *Philosophical Perspectives* 16: 105-34.

Bilgrami, A. 1992. *Belief and Meaning. The Unity and Locality of Mental Content*. Blackwell.

Black, T. 2008. Defending a Sensitive Neo-Moorean Invariantism. In Hendricks, V. (ed.), *New Waves in Epistemology*. Palgrave Macmillan.

Block, N. 1995. An Argument for Holism. *Proceedings of the Aristotelian Society*, 95: 151-169.

Borg, E. 2004. *Minimal Semantics*. Clarendon Press, Oxford.

Borg, E. 2009. Must a semantic minimalist be a semantic internalist?. *Proceedings of the Aristotelian Society Supplementary Volume.* LXXXIII: 31-51.

Borg, E. 2012. Semantics without pragmatics. In Allen, K. and Jaszczolt, K. (eds.). *The Cambridge Handbook of Pragmatics*. Cambridge: CUP: 513-528.

Brandom, R. 1994. *Making it Explicit*. Harvard University Press.

Breheny , R . 2003. A lexical account of implicit (bound) contextual dependence. In *Semantics and Linguistic Theory (SALT)* 13. Ithaca, NY: CLC Publications.

Bretherton, I. 1991. Intentional communication and the development of an understanding of mind. In Frye, D. & Moore, C. (Eds.), *Children's theories of mind: Mental states and social understanding*: Lawrence Erlbaum: 49-75.

Brown, J. 2005. Adapt or Die: The Death of Invariantism? *Philosophical Quarterly* 55 (219):263-285.

Burge, T. 1993. Content Preservation. *The Philosophical Review*, Vol. 102, 4: 457-488.

Cappelen, H. 2008. Content Relativism and Semantic Blindness. In G. Carpintero, M. & Koelbel, M. (eds.), *Relative Truth*. Oxford University Press.

Cappelen, H., Hawthorne, J. 2007 Locations and Binding, *Analysis* 67, 95-105.

Cappelen, H., Hawthorne. J. 2009. *Relativism and Monadic Truth*, Oxford University Press.

Cappelen, H., Lepore, E. 1997. On an Alleged Connection Between Indirect Quotation and Semantic Theory. *Mind and Language* 12: 278 –296.

Cappelen, H., Lepore, E. 2002. Indexicality, Binding, Anaphora and a Priori Truth. *Analysis* 62 (276): 271-81.

Cappelen, H., Lepore, E. 2005. *Insensitive Semantics: a defense of semantic minimalism and speech act pluralism*. Blackwell, Oxford.

Carpintero, M. 2001. Gricean Rational Reconstructions And The Semantics/Pragmatics Distinction. *Synthèse* 128, 1-2: 93-131.

Carston, R. 1988. Implicature, Explicature, and Truth-Theoretic Semantics. In Kempson, R. (ed.), *Mental Representations: the interface between language and* reality:155- 81. Cambridge University Press.

Carston, R. 1999. The Semantics-Pragmatics Distinction: a View from Relevance Theory. In Turner, K. (ed.) *The Semantics/Pragmatics Interface from Different Points of View*, Elsevier Science.

Carston, R. 2002. *Thoughts and Utterances: the pragmatics of explicit communication*. Blackwell, Oxford.

Carston, R. 2002a. Linguistic Meaning, Communicated Meaning and Cognitive Pragmatics. *Mind & Language*, 17: 127-148.

Carston, R. 2004. Explicature and semantics. In: S. Davis & B. Gillon (eds.) Semantics: A Reader. Oxford: Oxford University Press.

Castañeda, H.-N. 1999. *The Phenomeno-Logic of the I: Essays on Self-Consciousness* (edited by J. Hart and T. Kapitan). Bloomington : Indiana University Press.

Chomsky, N. 1957. *Syntactic Structures*, The Hague, Mouton.

Cohen, B., Murphy, G. 1984. Models of Concepts. *Cognitive Science* 8, 27-58.

Cohen, J., Rickless, S. 2007 Binding Arguments and Hidden Variables, *Analysis*, 67, 65-71.

Cohen, S. 1999. Contextualism, Skepticism, and the Structure of Reasons, *Philosophical Perspectives*, 13: 57–89.

Cohen, S. 2000, Contextualism and Skepticism, *Philosophical Topics* 10, pp. 94-107.

Cohen, S.: 1986, Knowledge and Context, *The Journal of Philosophy* 83: 574–583.

Cohen, S.: 1987, Knowledge, Context, and Social Standards, *Synthese* 73: 3–26.

Coliva, A., Moruzzi, S. Is Relativism Coherent?, MS.

Coliva, A., Moruzzi, S. Lost in disagreement, MSa.

Collins, J. 2007. Syntax, More or Less, *Mind* 116 (464): 805-850.

Corazza, E. 2007. Thinking the Unthinkable. In O'Rourke, M. & Washington, C. (ed.), *Situating Semantics: Essays on the Philosophy of John Perry, Cambridge*, MA: MIT Press.

Corazza, E., Dokic, J., 2007. Sense and Insensibility, or where Minimalism Meets Contextualism, in Preyer, G., Peter, G. (eds.). *Context-Sensitivity and Semantic Minimalism*, New York: Oxford University Press.

Corazza, E., Dokic, J., 2010. Situated Minimalism, *Synthèse* (online pre-publication).

Davidson, D. 1967 The Logical Form of Action Sentences, in N. Rescher (ed.), *The Logic of Decision and Action, Pittsburgh*: University of Pittsburgh Press.

Davidson, D. 1967a. Truth and Meaning, *Synthese*, 17: 304–23.

Davidson, D. 1973, Radical Interpretation, *Dialectica*, 27: 314–28.

Davidson, D. 1974, On the Very Idea of a Conceptual Scheme, *Proceedings and Addresses of the American Philosophical Association*, 47: 5–20.

Davidson, D. 1976, Reply to Foster, in G. Evans and J. McDowell eds. *Truth and Meaning: Essays in Semantics*. Oxford: Clarendon Press: 33–41.

Davidson, D. 1986/2005, A Nice Derangement of Epitaphs, reprinted in Davidson, D. 2005. *Truth, Language and History: Philosophical Essays*, with Introduction by M. Cavell, Oxford: Clarendon Press.

Davies, M. 1987. Tacit Knowledge and Semantic Theory: can a 5% Difference Matter?, *Mind* 96: 441-62.

DeRose, K. 1992. Contextualism and Knowledge Attributions, *Philosophy and Phenomenological Research*, 52/4: 913–29.

DeRose, K. 1995. Solving the Skeptical Problem. *Philosophical Review* 104: 1-52.

Devitt, M. (forthcoming) The Role of Intuitions, in Russell G., Graff-Fara D. (eds.) *A Companion to the Philosophy of Language*, Routledge.

Dummett, M. 1973. *Frege: Philosophy of Language*. London: Duckworth.

Dummett, M. 1976. What is a Theory of Meaning? (II) in G. Evans and J. McDowell eds. *Truth and Meaning: Essays in Semantics*. Oxford: Clarendon Press: 67–137.

Dummett, M. 1991. *The Logical Basis of Metaphysics*. Harvard Univ. Press.

Egan, A. 2007. Epistemic Modals, Relativism and Assertion, *Philosophical Studies* 133, pp. 1-22.

Egan, A. 2011. Relativism about Epistemic Modals, in Hales, S. (ed.) *A Companion to Relativism (Blackwell Companions to Philosophy)*, Blackwell.

Egan, A., Hawthorne, J., Weatherson, B. 2005. Epistemic Modals in Context. In Preyer, G., Peter, G. (eds.) *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press: 131-69.

Elugardo, R., Stainton, R. J. 2004. Shorthand, Syntactic Ellipsis, and the Pragmatic Determinants of What is Said. *Mind and Language* 19 (4):442–471.

Enc, M. 1987. Anchoring Conditions for Tense. *Linguistic Inquiry* 18: 633-57.

Evans, G. 1981. Semantic Theory and Tacit Knowledge, in Holtzmann, S. and Leich C. (eds.), *Wittgenstein: to Follow a Rule*, Routledge: 118-37.

Evans, G. 1982. *The Varieties of Reference*, OUP.

Field, H. 1977. Logic, Meaning and Conceptual Role. *The Journal of Philosophy* 74: 379-409.

Fodor, J. A. 1983, *The modularity of mind*. MIT Press, Cambridge, MA.

Foster, J. A. 1976 Meaning and Truth Theory, in G. Evans and J. McDowell eds. *Truth and Meaning: Essays in Semantics*. Oxford: Clarendon Press: 1–32.

Frege, G. 1892. Über Begriff und Gegenstand', in *Vierteljahresschrift für wissenschaftliche Philosophie*, 16: 192-205. Translated as *Concept and Object* by P. Geach, in Geach, P. and Black, M. (eds. and trans.), 1980. *Translations from the Philosophical Writings of Gottlob Frege*, Oxford: Blackwell;

Frege, G. 1893/1903. *Grundgesetze der Arithmetik*, Jena: Verlag Hermann Pohle, Volumes I-II.

Frege, G. 1956. The Thought, *Mind* , New Series, Vol. 65, No. 259: 289-311.

Fricker, E. 1995, Critical Notice: Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony, *Mind*, 104 : 393–411.

Gauker, C. 1997. Domain of Discourse, *Mind*, 106: 1–32.

Gauker, C. 2008. Zero Tolerance for Pragmatics. *Synthese* 165 (3):359–371.

Gauker, C. 2010. What Tipper is Ready For: A Semantics for Incomplete Predicates. *Noûs*, 46 (1):61-85.

Glanzberg 2007, Context, Content, and Relativism. *Philosophical Studies* 136 (1):1–29.

Grice, H. P. 1967/1989: Logic and Conversation, in Grice, P. 1989. *Studies in the Way of Words*, Cambridge (MA)-London: Harvard University Press.

Grice, H. P. 1969/1989: Utterer's Meaning and Intentions, in Grice (1989).

Grice, H. P. 1989: *Studies in the Ways of Words*. Cambridge, Mass.: Harvard University Press;

Grice, P. 1957. Meaning, *The Philosophical Review*, 66, 3: 377-88.

Hall, A. 2008. Free Enrichment or Hidden Indexicals? *Mind and Language* 23 (4):426-456.

Harman, G. 1972 Deep Structure as Logical Form', in Davidson, D. and Harman, G.(eds.), *Semantics of Natural Language*, Reidel, Dordrecht.

Harman, G. 1978. What is moral relativism? in A. I. Goldman and J. Kim (eds.), *Values and Morals* (Dordrecht, Holland; D. Reidel) pp. 143-161.

Harman, G., 1965, The Inference to the Best Explanation., *Philosophical Review* 74: 88–95.

Hawthorne, J. 2004. *Knowledge and lotteries*. Oxford: Oxford University Press.

Hawthorne, J. 2006. Testing for Context-Sensitivity, *Philosophy and Phenomenological Research*, Vol. LXXLII, No. 2.

Heim, I. & Kratzer, A. 1998. *Semantics in Generative Grammar*. Blackwell Publishing.

Horn, L. 2003. Pragmatics: Implicature, in Frawley, W.J. (ed.), *International Encyclopedia of Linguistics*, 2nd Ed., Oxford: Oxford University Press: 381–4.

Kaplan, D. 1989. Demonstratives, in Almog, J., Perry, J., Wettstein, H. (eds.): *Themes from Kaplan*; OUP.

Karttunen, L., and Peters, S. 1979. Conversational implicature, in Oh, C. & Dinneen, D. (edd.), *Syntax and Semantics, 11: Presupposition*: 1–56. New York: Academic Press.

Katz, J. 1972. *Semantic theory*. New York: Harper & Row.

Kempson, R. 1988. *Mental Representations: the interface between language and reality*, Cambridge University Press.

Kennedy, C. & McNally, L. 2010. Color, Context, and Compositionality. *Synthese* 174 (1).

King, J. 2003. Tense, Modality, and Semantic Values. *Philosophical Perspectives* 17 (1):195–246.

King, J., Stanley, J. 2005. Semantics, Pragmatics, and the Role of Semantic Content in Szabó Z.G. (ed.) *Semantics vs. Pragmatics*, Oxford University Press.

King, J. and Stanley, J. 2005. Semantics, Pragmatics, and the Role of Semantic Content, in Z. Szabó (ed.), *Semantics Versus Pragmatics* (Oxford: Oxford University Press): 111-164.

Kissine, M. 2010. From Contexts to Circumstances of Evaluation: is the Trade-off always Innocuous?, *Synthese.*

Kölbel, M. 2002. *Truth without Objectivity*, Routledge.

Kölbel, M. 2004a. Faultless disagreement. *Proceedings of the Aristotelian Society*, 104(1), 53–73.

Kölbel, M. 2004b. Indexical Relativism vs Genuine Relativism, *International Journal of Philosophical Studies* 12: 297–313.

Kompa, N. 2002. The Context-Sensitivity of Knowledge Ascriptions, *Grazer Philosophische Studien* 64: 1–18.

Korta, K., Perry, J., 2007, Radical minimalism, moderate contextualism. In Gerhard Preyer and Georg Peter (eds.), *Content and Context. Essays on Semantics and Pragmatics*, Oxford: Oxford University Press: 94–111.

Kripke, S. 1977. Speaker Reference and Semantic Reference, *Midwest Studies in Philosophy* II: 255-76.

Kripke, S. 1982. *Wittgenstein on Rules and Private Language*, Harvard: Harvard University Press.

Lahav, R. 1989. Against Compositionality: The Case of Adjectives. *Philosophical Studies*, 57: 261-279.

Lasersohn, P. 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28(4), 643–686.

Lasersohn, P. 2011. Context, relevant parts and (lack of) disagreement over taste. *Philosophical Studies*, 156: 433–439.

Leslie, S. J., 2007. Moderately Sensitive Semantics, in Preyer, G., Peter, G. (eds.). *Context-Sensitivity and Semantic Minimalism*, New York: Oxford University Press.

Levinson, S. 1983 *Pragmatics*. Cambridge University Press.

Levinson, S. 2000. *Presumptive Meanings: The Theory of Generalized Conversational Implicatures. Cambridge*: MIT Press.

Lewis, D. 1970. General Semantics, *Synthese*, 22: 18–67.

Lewis, D. 1979. Attitudes De Dicto and De Se. *Philosophical Review* 88 : 513-43, reprinted (with a postscript) in Lewis, D. 1983. *Philosophical Papers, Volume 1*. New York: Oxford University Press, pp. 133-159.

Lewis, D. 1980. Index, Context, and Content, in Kanger, S. and Öhman, S. (eds.), *Philosophy and Grammar*, Dordrecht: Reidel: 79–100; reprinted in Lewis, D. 1998. *Papers in Philosophical Logic*. Cambridge: Cambridge University Press.

Lewis, D. 1983. Introduction, *Philosophical Papers*, Volume I, Oxford: Oxford University Press.

Lewis, D. 1996, Elusive Knowledge, *Australasian Journal of Philosophy*, 74: reprinted in *Papers in Metaphysics and Epistemology*: 549–67.

Longworth, G. 2008. Linguistic Understanding and Knowledge, *Nous*, 42: 50–79.

Lopez de Sa, D. 2008. Presuppositions of Commonality. In Garcia-Carpèintero, M., Kolbel, M. (eds.), *Relative Truth*, Oxford University Press.

Ludlow, P. 2005. Contextualism and the New Linguistic Turn in Epistemology, in Preyer, G., Peter, G. (eds.) *Contextualism in Philosophy: Knowledge, Meaning and Truth*, OUP.

Lyons, J., 1997, Testimony, Induction, and Folk Psychology, *Australasian Journal of Philosophy*, 75:163–78.

MacFarlane, J. 2005a. Making Sense of Relative Truth, *Proceedings of the Aristotelian Society* 105, 321–39.

MacFarlane, J. 2005b. The Assessment Sensitivity of Knowledge Attributions, *Oxford Studies in Epistemology* 1: 197–233.

MacFarlane, J. 2007a. Semantic Minimalism and Nonindexical Contextualism, in *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics*, ed. G. Preyer and G. Peter (Oxford: Oxford University Press): 240–50.

MacFarlane, J. 2007b. Relativism and Disagreement, *Philosophical Studies* 132, 17–31.

MacFarlane, J. 2009. Nonindexical Contextualism, *Synthese* 166, 231–250.

MacFarlane, J. 2010. *Assessment Sensitivity*, MS.

Martí, L. 2006. Unarticulated Constituents Revisited. *Linguistics and Philosophy* 29 (2).

McConnell-Ginet, S. 1982. Adverbs and Logical Form, *Language* 58: 144-84.

McDowell, J. 1980 Meaning, Communication, and Knowledge, in van Straaten, Z. ed. *Philosophical Subjects: Essays Presented to P.F. Strawson*. Oxford: Oxford University Press: 117–39. Reprinted in McDowell, 1998: 29–50;

McDowell, J. 1993. Knowledge by Hearsay, in B. K. Matilal and A. Chakrabarti, eds. *Knowing from Words* (Kluwer, Dordrecht, Synthese Library vol. 230): 195–224.

Miller, A. 1997. Tacit Knowledge in Hale, B. and Wright, C. (eds.) *A Companion to the Philosophy of Language*, Blackwell: 146-74.

Millikan, R. G. 2004. *Varieties of Meaning: The Jean-Nicod lectures 2002*. MIT Press.

Millikan, R.G. 1984. *Language, Thought, and other Biological Categories: New Foundations for Realism*. The MIT Press.

Montague, R. 1974, Pragmatics in Montague, R. (ed.) *Formal Philosophy*: Yale University Press.

Murphy, G. 2002. *The Big Book of Concepts*, Cambridge, MA: MIT Press.

Neale, S 1990. *Descriptions* (Cambridge, Mass.: MIT Press).

Neale, S. 2004. This, That and the Other, in Reimer, M. and Bezuidenhout, A. (eds.), *Descriptions and Beyond*, OUP.

Neale, S. 2007. On Location. In O'Rourke, M. & Washington, C. (ed.), *Situating Semantics: Essays on the Philosophy of John Perry, Cambridge*, MA: MIT Press: 251–393.

Nunberg, J. 1995. Transfers of Meaning. *Journal of Semantics*.

Ogihara, T. 1996. *Tense, Attitudes, and Scope*. Dordrecht: Kluwer.

Pagin, P. 2005. Compositionality and context, in Preyer and G., Peter, G. (eds.), *Contextualism in Philosophy*, Oxford University Press: 303-48.

Pagin, P., Pelletier, J. 2007. Content, context and composition, in Peter, G., Preyer, G. (eds), *Content and Context. Essays on Semantics and Pragmatics*: 25-62, Oxford University Press.

Pagin, P., Westerståhl, D. 2010. Compositionality I: Definitions and Variants, *Philosophy Compass* 5:265-82.

Parsons, T. 1990. *Events in the Semantics of English. A Study in Subatomic Semantics*, Cambridge and London: MIT Press.

Partee, B.: 1989. Binding Implicit Variables in Quantified Contexts, *Proceedings of the Chicago Linguistics Society* 25, University of Chicago Press, Chicago: 342–65.

Pelletier, F. J. 1994. The Principle of Semantic Compositionality. *Topoi* 13: 11-24.

Pelletier, F. J. 2003, Context-dependence and Compositionality. *Mind and Language* 18: 148-161.

Perry, J. 1977. Frege on Demonstratives. *Philosophical Review* 86 (4):474-497.

Perry, J. 1979. The Problem of the Essential Indexical. *Nous* 13(1):3-21.

Perry, J. 1986. Thought without Representation, *Supplementary Proceedings of the Aristotelian Society*, 60: 137-52.

Perry, J. 1993. *The Problem of the Essential Indexical and Other Essays*. New York: Oxford University Press.

Perry, J. 2001. *Reference and Reflexivity*. Stanford, CA: CSLI Publications.

Predelli, S. 2005a. *Contexts: Meaning, Truth, and the Use of Language*: Oxford University Press.

Predelli, S. 2005b. Painted Leaves, Context, and Semantic Analysis. *Linguistics and Philosophy* 28.3: 351-74.

Prior, A.N. 1968. Now, *Noûs* 2.

Pupa, F., Troseth, E. 2011. Syntax and Interpretation, *Mind and Language*, Vol. 26, No. 2: 185–209.

Putnam, H. 1978. Realism and Reason, in *Meaning and the Moral Sciences*. London: Routledge and Kegan Paul, 1978: 123-138.

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

Quine Wv.O. 1968. Ontological Relativity, *Journal of Philosophy*, 65:185-212. reprinted in *Ontological Relativity and Other Essays*, New York: Columbia University Press, 1969: 26-68.

Quine, W. V. O. 1940. *Mathematical Logic*. Hrvard University Press.

Quine, W.V.O. 1953. Two Dogmas of Empiricism. in *From a Logical Point of View*, Harvard Univ. Press: 20-26.

Radford, A. 1981. *Transformational Syntax*, Cambridge University Press.

Radford, A. 1988. *Transformational Grammar*, Cambridge University Press.

Recanati, F. 1989. The Pragmatics of What is Said, *Mind and Language*, 4: 295-329.

Recanati, F. 1993, *Direct Reference: from Language to Thought*, Blackwell, Oxford.

Recanati, F. 2001. What is Said, *Synthese* 128: 75-91.

Recanati, F. 2002a. Unarticulated Constituents, *Linguistics and Philosophy* 25: 299-345.

Recanati, F. 2002b. Does Linguistic Communication Rest on Inference? *Mind & Language*, 17: 105-126.

Recanati, F. 2004. *Literal Meaning*. Cambridge Univ. Press, Cambridge.

Recanati, F. 2006. Literalism and Contextualism: Some Varieties. in Preyer, G. and Peter, G. (eds.). *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press.

Recanati, F. 2007. *Perspectival Thought*. OUP.

Recanati, F. 2010. *Truth-conditional Pragmatics*, Oxford University Press.

Richard, M. 2011. Relativistic content and disagreement, *Philosophical Studies* 156: 421–431.

Richard, M. 2008. *When Truth Gives Out*. Oxford University Press.

Richard, M. 2004. Contextualism and Relativism. *Philosophical Studies* 119 (1-2):215-242.

Rieppel, M. 2011. Stoic Disagreement and Belief Retention. *Pacific Philosophical Quarterly* 92:

243–262.

Rothschild, D. and Segal, G. 2009. Indexical Predicates. *Mind and Language* 24 (4):467-493.

Rumelhart, D. E. and Ortony, A. 1977. The representation of knowledge in memory. In Anderson, R. C., Spiro R. J. and Montague, W. E. (eds.). *Schooling in the Acquisition of Knowledge*. Hillsdale, NJ: Erlbaum.

Rumelhart, D.E. 1980. Schemata: the building blocks of cognition. In Spiro, R. J. et al. (eds) *Theoretical Issues in Reading Comprehension*, Hillsdale, NJ: Lawrence Erlbaum.

Rumfitt, I. 2005 Meaning and Understanding, in F. Jackson and M. Smith eds. *The Oxford Companion to Contemporary Philosophy*. Oxford: OUP: 427– 423.

Rysiew, P., 2000, Testimony, Simulation, and the Limits of Induction. *Australasian Journal of Philosophy*, 78: 269–274.

Searle, J. 1969. *Speech Acts: An essay in the philosophy of language*, Cambridge: Cambridge University Press.

Searle, J. 1975. A taxonomy of illocutionary acts, in Gunderson, K. (ed.), *Language. Mind and Knowledge: Minnesota Studies in the Philosophy of Science*, vol. VII, Minnesota: University of Minnesota Press: 344–69.

Searle, J. 1980. The Background of Meaning in J. Searle, F. Kiefer & M. Bierwisch (eds.) *Speech Act Theory and Pragmatics*, Riedel.

Sellars, W. 1954. Presupposing. *The Philosophical Review*. 63, 2 (Apr., 1954): 197-215.

Sennet, A. 2008 The Binding Argument and Pragmatic Enrichment, or, Why Philosophers Care Even More than Weathermen about 'Raining', *Philosophy Compass* 3/1, 135–157.

Soames, S. 2002 *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. New York: Oxford University Press.

Sosa E., Kim, J. (eds.) 2000, *Epistemology: An Anthology*, Malden, MA: Blackwell: 340-53.

Sperber, D. 1995. How do we communicate? In Brockman, J. & Matson, K. (eds) *How things are: A science toolkit for the mind*. New York : Morrow: 191-199.

Sperber, D. 2000. Metarepresentation in an evolutionary perspective. in Sperber, D. (ed.) *Metarepresentation: A Multi-Disciplinary Perspective*. Oxford University Press.

Sperber, D., Wilson, D. 1986/1995: *Relevance*. Cambridge: Harvard University Press.

Sperber, D., Wilson, D. 1998. The Mapping between the Mental and Public Lexicon. In Carruthers, P. and Bouchers, J. (eds.), *Language and Thought: Interdisciplinary Themes*. Cambridge Univ. Press: 184-200.

Sperber, D., Wilson, D. 2002. Pragmatics, Modularity and Mind-reading. In *Mind and Language*: 17. 3-23.

Sperber, D. 1994. Understanding verbal understanding. In Khalfa, J. (ed.) *What is Intelligence?* Cambridge University Press: 179-198.

Stainton, R. J. 2005. In Defense of Non-Sentential Assertions. In Szabó, Z. G. (ed.), Semantics Versus Pragmatics. Oxford University Press.

Stanley, J. 2000. Context and Logical Form, *Linguistics and Philosophy* 23: 391-424.

Stanley, J. 2002. Making it Articulated. *Mind & Language*, Vol. 17:149–168.

Stanley, J. 2005a, *Knowledge and Practical Interests*, (Oxford: Oxford University Press).

Stanley, J. 2005b. Semantics in Context. In Preyer, G. & Peter, G. (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford University Press.

Stanley, J., Szabó, Z. 2000. On Quantifier Domain Restriction. *Mind and Language* 15: 219-61.

Stanley, J., Williamson, T. 1995, Quantifiers and Context-Dependence, *Analysis*, 55: 291–5.

Stephenson, T. 2007. Judge Dependence, epistemic modals and predicates of personal taste, *Linguistics and Philosophy*, 30: 487-525.

Stojanovic, I. 2007. Talking about taste: disagreement, implicit arguments, and relative truth. *Linguistic and Philosophy* 30: 691–706.

Sundell, T. 2011. Disagreements About Taste. *Philosophical Studies* 155 (2):267-288.

Szabó, Z. G. 2005. *Semantics Vs. Pragmatics*. Oxford University Press.

Szabó, Z.G. 2001. Adjectives in Context. In Harrish R. & Kenesei, I. (eds.), *Perspectives on Semantics, Pragmatics, and Discourse*. John Benjamins Publishing Company.

Tarski, A. 1936. Der Wahrheitsbegriff in den formalisierten Sprachen. *Studia Philosophica*, 1: 261–405.

Taylor, K. 2001. Sex, Breakfast, and Descriptus Interruptus, *Synthese* 128: 45–61.

Travis, C. 1975. *Saying and Understanding*, Oxford, Blackwell.

Travis, C. 1985/2008. On What is Strictly Speaking True, *Canadian Journal of Philosophy* 15: 187-229, reprinted in Travis (2008).

Travis, C. 1996. Meaning's Role in Truth, *Mind* 100: 451-66.

Travis, C. 1997: Pragmatics, in Hale, B., Wright, C. (eds.): *A Companion to the Philosophy of Language*, Oxford: Blackwell.

Travis, C. 2008. *Occasion Sensitivity*, New York: Oxford University Press.

Turner, K. (ed.) 1999. *The Semantics/Pragmatics Interface from Different Points of View*, Elsevier Science.

Turri, J. 2010. Epistemic Invariantism and Speech Act Contextualism. Philosophical Review 119 (1): 77-95.

Unger, P.: 1984, *Philosophical Relativity*, Basil Blackwell, Oxford.

Unger, P.: 1986, The Cone Model of Knowledge, *Philosophical Topics* 14: 125–178.

Vendler, Z. 1957. Verbs and times. *Philosophical Review* 56: 143-160.

von Fintel, K., Gillies, A. 2008. Cia Leaks. *Philosophical Review* 117 (1):77-98.

Wilson, D. 2000. Metarepresentation in Linguistic Communication. In Sperber, D. (ed.) *Metarepresentation: A Multi-Disciplinary Perspective*. Oxford University Press.

Wittgenstein, L. 1953. *Philosophical Investigations*, Anscombe, G.E.M., Rhees, R. (eds.), Anscombe, G.E.M. (trans.), Oxford: Blackwell.

Wittgenstein, L. 1922. *Tractatus Logico-Philosophicus*, Ogden, C.K. (trans.), London: Routledge & Kegan Paul.

Wright, 2001. On Being a Quandary. *Mind* 110, No. 437 (Jan., 2001): 45-98.

Wright, C. 2006. Intuitionism, Realism, Relativism and Rhubarb. in Greenough, P. and Lynch M. (eds.), *Truth and Realism*, Oxford University Press.

Zeman, D. 2011. *The Semantics of Implicit Content*, Ph.D Thesis, University of Barcelona.