What Is Said. An Inquiry into Reference, Meaning, and Content.

Isidora Stojanovic

▶ To cite this version:

HAL Id: tel-00144054
https://theses.hal.science/tel-00144054

Submitted on 30 Apr 2007

WHAT IS SAID

AN INQUIRY INTO REFERENCE, MEANING AND CONTENT

Isidora Stojanovic

Abstract

I investigate the relationship among the notions of meaning, content, and *what is said*. It is widely held that *indexicals* – words like 'this', 'I', or 'today' – contribute their reference, and nothing but their reference, to the semantic content, and thereby undermine any tentative identification of semantic content with lexical meaning. Against the mainstream view, I argue that semantic content *is* lexical meaning, for indexical and non-indexical expressions alike. In Chapter 1, I lay down this proposal in all due detail, explaining how to think of the semantic content of sentences containing indexicals, and articulating the relationship between content, truth, and reference. In Chapter 2, I present a number of problems for the existing accounts of *what is said*, and then show that if we think of semantic content along the lines of my proposal, we may account for the problematic cases while identifying the asserted content (or what is said) with semantic content. In Chapter 3, I show how my account extends, on the one hand, to definite descriptions and proper names, and, on the other, to epistemic modals and predicates of taste.

Preface

The thesis that I defend in this work is that the semantic content of an indexical expression, such as 'this' or 'I', is nothing more or less than its lexical meaning. This is a controversial thesis. The mainstream view, inherited from David Kaplan, is that indexicals contribute their reference to the semantic content. In the view that I will defend, "the reference of an indexical" becomes a misleading expression. I agree that direct reference is an important phenomenon, and that the things referred to matter to the truth value, but it is the speaker who is referring to things directly, rather than the words used by the speaker.

At this stage, I merely want to forestall a few possible confusions, and then announce what will happen in the following chapters. My view departs from the mainstream view in that it holds that the semantic content of indexicals is their lexical meaning. But apart from that, there is little divergence between the two views. In particular, I follow the tradition of truth-conditional, model-theoretic approach to meaning, initiated by Richard Montague and well-established in contemporary formal semantics. My proposal does not require any revision of the major assumptions regarding natural language semantics.

It may be said that my view on indexicals radically departs from the mainstream view, while being at the same time extremely conservative. This is because the view put forward by Kaplan himself in the 70's was a radical departure from what was then the mainstream view, namely the index-theory of David Lewis, Hans Kamp, Richard Montague, and others. Kaplan thought that indexicals posed a major obstacle to that theory and that they called for a "two-tiered" theory of meaning. My view rejects Kaplan's distinction between characters and contents by holding that the lexical meaning (which Kaplan

tried to capture with his notion of character) just *is* the semantic content. In this sense, this work may be seen as an attempt to rehabilitate the pre-Kaplanian truth-conditional semantics.

It might be a good idea to state at the outset a number points on which I believe my view to be superior to the standard view. I will go from the least to the most important.

- The view that I am proposing is more economical, and presumably more elegant, than the mainstream view. There is no distinction between "two kinds of meaning" such as characters vs. contents. There is only one level of meaning, and that is, invariably, the lexical meaning or the semantic content.

- I can provide a neat definition of the semantics/pragmatics distinction. Semantic content is the lexically encoded, truth-conditional, context-invariant content, while anything that depends on the context (including reference) belongs to the realm of pragmatics. (cf. 1.4).

- I can purport to provide a unified account of deixis and anaphora: the meaning (or content) of a pronoun like 'she' remains constant, regardless of whether the pronoun is used deictically or anaphorically. (cf. 1.8).

- The view that I am proposing may be nicely extended to expressions that share some important features with indexicals (while being interestingly different) and that have received considerable attention in the past couple of years, such as epistemic modals and predicates of taste. (cf. 3.2-3.3).

- The proposed notion of semantic content may stand for the asserted content. The theory of what is said put forward in Chapter 2 (what is said = lexical meaning = semantic content) is superior not only to Kaplan's theory of what is said, but also to the other theories currently represented.

Now that I have announced some desirable consequences of the view that I will defend, let me briefly say how I will proceed. In Chapter 1, I lay down and defend the idea that semantic content is lexical meaning. In doing so, I will clarify how we are to think of lexical meaning, and I will explain where

and how reference comes into the picture. I will also discuss the interaction between indexicals and intensional operators, the role of double indexing, and certain related issues and phenomena, such as presupposition and anaphora. I will try to present my view in an informal and accessible way, confining the formal account to the Appendix.

Chapter 2 is where I lay down a series of problems for Kaplan's and other accounts of *what is said*. At a first glance, one might think that the problems pointed out require introducing further "tiers" of meaning, or a plurality of contents. In an even more pessimistic vein, one might fear that the notion of what is said is just too vague and versatile to be of interest to semantics. However, I will show that the cases problematic for the existing accounts of what is said are no longer problems once we start to think of semantic content along the lines of Chapter 1. In my view, then, what is said is the semantic content, and the latter is the lexical meaning. However, we will also have an independently motivated notion of *subject matter*, or that of which what-is-said is said, and I will show how the two notions (viz. what is said, and that of which it is said), together with some assumptions on our linguistic practices of reporting what is said, make it possible to account both for the cases that had originally motivated Kaplan's theory, and for those that put it into jeopardy.

Finally, Chapter 3 will take us beyond indexicality to, on the one hand, definite descriptions and proper names, which, together with indexicals, are believed to form the class of singular terms, and, on the other, to epistemic modals and predicates of taste, which are context-sensitive, yet significantly different from indexicals. Of course, any thorough discussion of these topics would go far beyond the scope of the present work. I will only discuss these topics to the extent that my account of indexicals has some interesting bearings on them, as in the case of epistemical modals and taste predicates, or because a theory of indexicals that says nothing about their similarities and differences with definite descriptions, on the one hand, and proper names, on the other, is, in some sense, incomplete.

Acknowledgments

This work would not have existed, at least not in its present form, if I had not been very lucky to have worked with some quite extraordinary people. The order of acknowledgments is chronological. The person who initiated me to philosophy of languge is François Recanati, and, as my former adviser at the Ecole Polytechnique and present colleague at the Jean Nicod Institute, he is someone to whom I owe a lot and am extremely thankful. I also wish to thank Joseph Almog and David Kaplan, who, back in 1998, made it possible for me to spend a quarter at UCLA, which had a great impact on my philosophical development. I also thank Philippe Schlenker, from whom I learnt a lot, and who has been a good friend, too. At Stanford, I owe very special thanks to my adviser John Perry, who gave me extensive feedback on earliers draft of this work, made many insightful suggestions, has always been a great listener, and, most importantly, a wonderful person. If this work had not been already rather long, I would have talked more about Perry's view and the respects in which the view defended here is similar to or differs from his. I also benefited greatly from discussions with Kenneth Taylor and Mark Crimmins. Then, I wish to thank John MacFarlane for detailed comments on an earlier draft. Finally, for over a year, I have had many opportunities to discuss with Stefano Predelli. Those discussions have had a great impact on my work. He has also been a very dear and helpful friend, and I wish to thank him for that, too.

Table of Contents

Chapter 1. **Indexicals Revisited**.

It is commonly held that what indexicals contribute to semantic content are their referents. My thesis is that this is wrong. I will defend a view on which semantic content just is lexical meaning, in the case of indexical and non-indexical expressions alike. In laying down this view, I will explain where and how reference comes into the picture.

Chapter 2. **What Is Said**

A major motivation for Kaplan's notion of content is that it should stand for the asserted content, or what is said. *In this chapter, I lay down a series of problems for the existing theories of what is said. However, if we think of semantic content along the lines of the theory presented in Chapter 1, then we may identify semantic content with what is said (and identify both with lexical meaning).*

Chapter 3. **Beyond Indexicals**

In this chapter, I show how the view laid down in the previous chapters may be extended to certain non-indexical expressions, and how it sheds light on a number of issues, some of which have been known for a long time, while others have been brought up only recently.

Chapter 1

Indexicals Revisited

It is commonly held that what indexicals contribute to semantic content are their referents. My thesis is that this is wrong. I will defend a view on which semantic content just *is* lexical meaning, in the case of indexical and non-indexical expressions alike. In laying down the view, I will explain where and how reference comes into the picture.

1.1. Introduction

The mainstream approach to indexicals has its origin in David Kaplan's pioneering work on what he then called *demonstratives* (Kaplan (1977)). Throughout this work, I will use the term 'indexical' in the general sense that includes not only "pure" indexicals such as 'I' and 'today', but also, for instance, the third person pronouns ('he', 'she') and demonstrative pronouns ('this', 'that').

Suppose that, pointing at Suleïka, I tell you:

(1) She is dangerous.

The mainstream, Kaplanian, or "direct referentialist" approach to indexicals may be characterized by the following claims:

(i) with my utterance of (1), I will communicate to you something about Suleïka herself, namely, that she is dangerous;

(ii) I can do so because the sentence uttered in (1), relative to its context of utterance, expresses the proposition that Suleïka is dangerous (at the time of utterance);

(iii) that same "singular" proposition is the semantic, or truth-conditional, content of the sentence uttered in (1) (relative to its context of utterance);

(iv) what the pronoun 'she' contributes to the semantic content of the sentence in (1) (relative to its context of utterance) is Suleïka herself, – and nothing else.

Inevitably, squeezing the mainstream view into these four claims requires a certain amount of simplifying. For one thing, in Kaplan's own (formal) account, the pronoun's contribution to the content is not Suleïka herself, but the "individual concept" of her, which is a function that for any point of evaluation returns Suleïka. Nevertheless, it has become customary to talk of the semantic contribution of indexicals, and directly referential expressions more generally, as being the referent itself. I shall follow the custom. As will become clear shortly, my concerns regarding the mainstream view do not at all hinge on the idea that an individual may be part of the semantic content, and will certainly not be remedied by trading individuals for individual concepts.

Let me say at the outset where and how the position that I am going to defend in this work departs from the mainstream view. First of all, the two views agree on point (i), that is, on the idea that in (1) I am talking about Suleïka, and that I may thereby communicate to you something *directly* about this very individual, namely, that she is dangerous. It may be acknowledged that agreement on this point shows that my position is not Fregean. Recall that for Frege, what is communicated is the *sense* (Sinn). An utterance of (1) communicates something "about" Suleïka only to the extent that she happens to be the *denotation* (Bedeutung) determined by some constituent sense (cf. Frege (1897)). On the other hand, the position endorsed here is closer to direct reference theories than to Frege, since I hold that what a persons says may be, and often is, *directly* about objects, people, places, and that reference to those

things need not be mediated by any sense-like entity. I am not going to argue for point (i). If the reader believes that my utterance is about Suleïka only to the extent that she is suitably related to some sense-like abstract object associated with my utterance, then let me ask you to take my talk about direct reference simply as a *façon de parler*. For, I believe that my main thesis, which concerns the indexicals' contribution to semantic content and to what is said, would still go through even in a Fregean framework, in which whenever we seem refer to things directly, it actually turns out that reference is mediated by some sort of sense or "mode of presentation".

Let us turn back to the mainstream referentialist view. The points on which that view and mine disagree are points (iii) and (iv). I will defend a position that says that Suleïka is not the semantic content of 'she', nor is she that by which the pronoun contributes to the semantic content of the sentence in (1). The pronoun's contribution is, I suggest, merely the condition of being female, which we take to be "lexically encoded" in the 3rd person pronoun 'she'. Furthermore, in my view, Suleïka is not at all part of the semantic content of (1), and that is why point (iii) must go as well.

Finally, what about point (ii)? I can neither say that I agree nor that I disagree. In the view that I will defend, it will certainly be possible to *define* what it takes for an utterance to "express a singular proposition", and to define it in such a way that (1) ends up expressing the singular proposition *that Suleïka is dangerous (at the time of (1))*. Suppose for a second that we have done so. Then, what is also true about my view is that the notion of proposition (be it singular or general), and the related notion of expressing, will prove to be utterly useless. First, as already said, such a singular proposition is not going to be the semantic content of the sentence uttered. Second, it will play no role in determining the semantic content of the sentence in (1) or of any other sentence – no role, in particular, in determining the semantic content of complex sentences such as "It is possible that she is dangerous" or "Tarek believes that she is dangerous." Thirdly, such a

"proposition expressed" is not going to be the asserted content, or *what is said* – in the account that I will propose, *semantic* content is the asserted content. Fourthly, it is not the truth-conditional content – again, semantic content is the truth-conditional content. So, in sum, since I fail to see any useful role for such a notion of "proposition expressed," from now on, I will avoid using the term 'proposition' (except in describing other people's views). But anyone who is fond of singular propositions and who thinks that they can be useful will be able to get them pretty much for free in the view that I will put forward.

1.2. Referring Directly vs. Using Indexicals

If I want to convince you that my view is more attractive than the mainstream view, the typical thing to do would be to assault the "enemy view" with a bunch of problems, show that it has no way to go, and then offer my own view as the best available alternative. Whatever the merits of this way proceeding, I would like to try something else. I would like to start by presenting my own view, and to do it as if it were "from scratch." So what I need to ask you is that, at least for this first chapter, you pretend to be a novice, as if this were the first time that you ever heard about indexicals and reference, and that you let yourself walk through the informal picture that I am about to give, looking at it with a fresh eye.[1]

1.2.1. Direct Reference "Par Excellence"
I take it to be something of a platitude that communication presupposes that we should be able to convey information *about* people and things around us, and that this, in turn, strongly suggests that we should be able to *refer* to those things *directly*. Here is an example of what I take to be a paradigmatic case of

---

[1] In Chapter 2, though, I will use the usual method: I will present problems for the standard, and the less standard, views on what is said, and then show how the problems dispel if one adopts the view that I am offering in alternative.

direct reference. Suppose that we are staring at a certain painting – let's say, at Rembrandt's masterpiece *The Night Watch*, and I simply say:

(1) Impressive!

Then I will be referring to that very painting of Rembrandt, and I will be saying *of it* that it is impressive.

Here is another example of what I like to think of as direct reference. Suppose that we are on the Main Quad at Stanford, it starts raining, and you say:

(2) It's raining!

The place where you say that it is raining, the Main Quad at Stanford, has been referred to directly. One key similarity between the two cases is that neither of us had to employ any expression for the thing and the place that we were talking about, let alone find some descriptive way of identifying them. Another key similarity is that the truth of our utterances depends on the things that we were talking about. Whether (1) is true depends on the facts that obtain of that painting of Rembrandt, and on whether it is indeed impressive.[2] Similarly, whether (2) is true depends on the facts (meteorological facts, if you wish) that obtain of that very place, the Main Quad at Stanford, at the time when (2) is uttered.

It is this form of reference, reference that enables us to talk about a particular thing even though we do not need to use any expression for it, that I consider to be direct reference *par excellence*. Reference supported by words, be they indexicals, proper names, or definite descriptions, is, in my view, merely parasitic on this other, more basic form of reference.

1.2.2. Direct Reference and the Use of Indexicals

It is easy to confuse direct reference with the use of indexicals. A possible explanation of why indexicals are so often thought of as devices of direct

---

[2] Let us pretend at this stage that there is actually a matter of fact as to whether a given thing is impressive or not. Predicates of taste (such as 'impressive') will come under scrutiny in 3.2.

reference is that in theorizing about them, philosophers have often focused on those uses of indexicals on which they do no interesting work – neither for semantic interpretation nor for communication – rather, they merely "articulate" the reference. Consider:

(3) This is impressive!

(4) It's raining here!

These appear to be pretty much equivalent ways of saying, or communicating, what was communicated in (1) and (2). And it is precisely cases such as (3), and perhaps (4),[3] that seem to support the idea that we use indexicals in order to refer to things, and that indexicals merely stand for their reference.

But that idea I believe to be wrong. Of course, if one focuses on those cases in which the speaker has used some indexical word when she could have just as well used no word at all, as in the pair (3)-(1), it becomes somewhat difficult to say what else the indexical is there for, other than to articulate the reference. In such cases, the speaker might have just as well referred to the relevant thing *directly*, making it clear to her audience what she was referring to, without the help of *any* linguistic expression. But the role of indexicals, when they are not "idle", is precisely to help the speaker refer to something, and to help the hearer figure out what it is to which the speaker has referred. So the picture is roughly as follows. Direct reference comes first, to put it so, and it does not require the use of any expression, indexical or not, that would "stand for" the thing referred to. Nevertheless, direct reference may be accompanied by the use of an indexical. When this happens, the speaker will typically use the indexical in order to help her audience figure out what it is to which *she*, *qua* speaker, is referring. In other words, you do not need indexicals in order to

---

[3] It is a fact, though, that (4) sounds quite odd in any situation in which the speaker could have just said "It's raining", as in the situation of (2). The indexical 'here' is normally used to mark a contrast, and if there is no contrast to be marked, it becomes redundant. Unlike the case of locational adverbs such as 'here', the use 'this' in (3) is justified from the standpoint of syntax alone. For discussion regarding sentences that lack mandatory grammatical subjects, see, for instance, Stainton (2000).

refer to things, but you might still need them in order to make it clear to the others what it is to which you are referring.[4] And the way in which indexicals help the hearer figure out what the speaker is referring to is by constraining the range of possible referents, by means of conditions lexically encoded in their meaning. The next section expands on this idea.

1.2.3. The Lexical Meaning of Indexicals

One of the greatest merits of Kaplan's work on indexicals was to show that indexicals have an interesting sort of contextually stable meaning, and to propose a model-theoretic analysis of it. But what is, informally, the lexical meaning of a word such as 'this'? It is, roughly, what competent language users know, and what dictionaries try to provide. For example, the meaning of 'this' encodes the information that the thing talked about is salient and proximal to the speaker, while the meaning of 'today' presumably bids the interpreter to identify the day on which the utterance is taking place.

Though it is convenient to talk of the meaning of 'this' as something like "the most salient thing that is proximal to the speaker", or of the meaning of 'I' as "the speaker", or of 'today' as "the day of utterance", let me make one important thing crystal-clear at the very outset. We should not expect to be able to pin down the lexical meaning of an indexical by means of some determinate definite description. We may want to use such descriptions simply as *approximations* to the lexical meaning of the word, just as we may use the description "hand-held instrument containing an interior strip of solid material that produces marks used to write or draw" as an approximation to the lexical meaning of the common noun 'pencil'. But of course, no one believes that this description *is* the lexical meaning of the word 'pencil'. There is no reason, then, to think that the condition of speakerhood, say, should capture the lexical meaning of the indexical 'I', or that salience should be an essential component to the lexical meaning of the demonstrative 'this', and so forth.

---

[4] John Perry has urged me to clarify this subtle point.

Let me now give you my own picture of how indexicals work of the role that their lexical meaning play. Suppose that we are at a party, and I want to warn you that Suleïka is dangerous. Suppose that it is mostly men at the party, and that there is this couple, a man and a woman, the man is Tarek and the woman is Suleïka, and both of them are very loud and noticeable, and I say:

(5) She is dangerous.

The 3rd person pronoun 'she' has only a very "poor" lexical meaning. All that is lexically encoded is that the person referred to should be *female*.[5] But even this very poor condition is doing something very useful in communication. It is helping you, *qua* hearer, to figure out that it is Suleïka to whom I am referring, since she is the only one, among the things or individuals to which, for all you know, I might be referring, who satisfies the condition associated with the pronoun 'she'.

To see that the pronoun in (5) is not idle, contrast that case with a similar situation in which Suleïka is extremely loud and noticeable and our attention is already focused on her, and I utter the same sentence as before, namely:

(6) She is dangerous.

In both this and the previous case, I, *qua* speaker, am directly referring to Suleïka, intending to communicate something about her. But in this case, unlike the previous one, I do not need to worry much about helping you figure out to whom I am referring. Suleïka is already so salient that she will be picked up by default as my intended reference. In this kind of case, the lexical meaning of the pronoun is, as I like to say, *idle*. The conditions lexically encoded in the pronoun are not put at work for communicational purposes.

---

[5] Some theoreticians hold that the negative conditions of not being the speaker and not being the addressee are also encoded in the lexical meaning of the 3rd person pronoun. My position is that those conditions are not part of the lexical meaning. I owe you, then, an explanation of why it is wrong for a female speaker, say Suleïka, to say 'She is happy' in order to express that she herself is happy, given that, being female and salient enough, Suleïka satisfies the conditions lexically associated with 'she'. I must ask you to wait until 1.9.2. for an explanation of why this is so.

The pronoun is mainly used for syntactic reasons. In a language in which the grammatical subject is not mandatory, I could have simply said:

(7) Dangerous.

Note that (7) would not have worked in the context of (5), because there, Tarek and Suleïka were equally salient and equally likely to be picked up as the default reference.

Finally, consider a situation pretty much like that of (5), except that the strange couple are now Suleïka and Maria, and are both equally salient and noticeable. I again want to inform you that Suleïka is dangerous, but now it will not work to simply say "She is dangerous." For, I have given you no hints whatsoever that will let you decide whether it is Suleïka or Maria that I am referring to. In such situations, there are typically two ways to go. One is to somehow enhance the salience of the person to which I am referring, for example, by pointing at her with my finger. Once I make sure that Suleïka has been raised in salience over Maria, I can go back and just use the pronoun 'she', as in (5) or (6). But in many contexts, that is not the best strategy. (For instance, it is not polite to point at people with your finger at parties, and even if I did, if Suleïka and Maria are side by side, it would still be difficult for you to see which of the two I am pointing at, and so on.) A better strategy is to use more language. For instance, I might say:

(8) That woman with a pink hat is dangerous.

The lexical meaning of the complex demonstrative in (8) provides you not only with the condition that the pronoun 'she' provides you with (being female, and perhaps being salient) but an additional condition ('with a pink hat'), which will help you decide that I must be referring to Suleïka.

1.3. Semantic Content as Lexical Meaning

My thesis is that the semantic content of any given sentence, whether or not it contains indexicals, simply is the lexically encoded content, or the lexical meaning of the sentence. But now, something needs to be said on how to think of this lexical meaning. For, it clearly cannot be the same thing as, say, Kaplan's *character*, because character is defined in terms of content: character is, recall, a function from contexts to contents. For us, too, lexical meaning is going to be a certain function. In the sentential case, it will be a function from certain things to truth values. But which things? That will be the question with which the next section deals.

1.3.1. A Truth-Conditional Approach to Lexical Meaning

Before we answer the question in full generality, let us look at a few particular cases. Let us start with the simple ones, those where the speaker is directly referring to something, and is saying something about that to which she is referring, and there are no indexicals at play. Thus recall our conversation in front of Rembrandt's *Night Watch*, where I say:

(1) Impressive!

What will be the semantic content associated with my utterance of (1)? We already know two things: that it is lexically encoded in the sentence uttered in (1), and that it will eventually return a truth value. Now, the sentence uttered reduces to the predicate 'impressive'. So the semantic content requires at least an object before it can yield a truth value. Furthermore, as we know from the work on modalities and tenses, the truth of natural language sentences systematically varies along these two dimensions: worlds and times.[6] The semantic content associated with (1) will be, then, conceived of as a function that takes an object, a time, and a possible world, and returns a truth value.

---

[6] The variability of truth across times is more controversial. Many philosophers wish to say that natural language sentences, even when they contain no overt temporal indexical, are still time-specific and anchored to the time of utterance (as if they were prefixed with 'now'). For further discussion of these issues, I must ask you to wait until section 2.4.

And this is pretty much the final account, with the only exception that we should not limit ourselves to a single object. In full generality, the semantic content of a sentence is a function that takes an ordered set of objects – itself modeled as a function that assigns values to variables – and it takes a time and a possible world, and returns a truth value.[7]

The semantic content of (1) does not involve the object to which I am referring and of which I am saying that it is impressive. Rembrandt's painting is not part of the semantic content, not any more than the actual world is. But observe that reference still interestingly connects with semantic content. For, in order to get to the actual truth *value* of (1), you must feed an object into the semantic content of the sentence in (1). The object fed in, at which you "evaluate" the sentence for its truth value, will normally be the object to which the speaker is referring while uttering the sentence.

Let us now turn to indexicals. Consider:

(2) This is impressive.

We need to add very little to the picture already set up for the case of (1). On a first approximation, the semantic content of (2) will again be a function that takes an object, a time and a possible world, and returns a truth value. For instance, if you give it Rembrandt's *Night Watch,* January 2007, and the actual world, we will expect it to return value True. On the other hand, if you give it something that is not impressive – let's say, an empty jar of peanut butter – and you give it the same time and world, it will return value False.[8]

I have been repeatedly stressing how indexicals contribute their lexical meanings to the semantic content. The difference between the contents of (1) and (2) will, then, amount to the following. If you give object o, time t and world w as inputs to the content of (2), it will return value True if, and only if,

---

[7] We will soon see that there is a need for double indexing, and in the version of double indexed semantics that we will opt for, semantic content will, technically, be a function that takes not just one, but two possible worlds, and not one, but two times.
[8] Just as a reminder, we are still pretending that there is a matter of fact as to whether something is impressive or not.

o is impressive at *t* and *w*, and (– here comes the indexical's contribution –) if *o* is "a this" at *t* and *w*, that is, if it is proximal to some speaker and salient to her audience (at *t* and *w*).

Consider another example:

(3) She is dangerous.

What is the semantic content of (3)? A function that (so far) takes a triple (*o, t, w*) and returns a truth value. Suppose that (at *t, w*) Suleïka and Tarek are dangerous individuals, and Maria isn't. The semantic content of (3) will return value True for Suleïka, and will return value False for Tarek and Maria. The reason why it is false for Tarek, who is dangerous, is that he is male. The content of (3) will only return value True if the input individual is dangerous *and female* (at *t, w*). The latter condition is what is contributed by the pronoun 'she', since that is the condition lexically encoded in the pronoun.

1.3.2. *Some Remarks On the Notion of Truth Conditions*

In philosophical literature, there seems to be some unclarity lingering over the term of truth conditions. The way in which this term as well as the term 'truth-conditional content' are used by some philosophers shows that they take it to be implicit in those terms that the conditions at stake need to be *worldly* conditions. In other words, for those philosophers, some condition, associated with some sentence S, counts as a truth condition only if it specifies *what the world must be like* for the sentence to be true.

I submit that there is no good theoretical or practical reason to insist that truth conditions should necessarily be *worldly* conditions. Maybe there is some historical reason. But once we see that the assumption is arbitrary, there is no reason to stick to it, not even for historical reasons. Instead, we should rehabilitate the term 'truth conditions' so that it means what it really means – and that is, simply, truth conditions, which may be conditions pretty much on anything. For example, take the condition of being bigger than. That is a condition on *pairs* of objects – a condition that obtains when the first one is

bigger than the second. Why not think of this condition as being a "truth condition" associated with the open sentence "x is bigger than y"? I cannot see any possible reason why not. Now, as I have already suggested, natural language sentences that contain indexicals are such "open" sentences. And what I have proposed as their semantic content – that is, functions from an ordered set of objects, a time and a world, to truth values – are, I submit, perfectly fine truth conditions. To be sure, they are not *worldly* conditions. But the assumption that all truth conditions should be conditions on what the world must be like is, again, arbitrary, and there is no reason to keep it. Once the confusion and unclarity have been lifted, we can appreciate the fact that the semantic contents proposed here are *truth-conditional* contents.

1.4. The Semantics/Pragmatics Distinction

The problem of the semantics/pragmatics distinction – by which I mean the problem of defining 'semantics' and 'pragmatics' in a coherent and intuitively plausible way – has come to considerable discussion in the past decade.[9] While it appears that there is indeed such a problem for the mainstream view, and that the view must eventually make some arbitrary decisions and give up some plausible criteria for what may count as 'semantics', as will be explained in what follows, it turns out that there are no such problems for the view defended here. By insisting that even in the case of "directly referential" expressions such as indexicals, reference is not part of the semantic content, and that all there is to the semantic content is the lexically encoded meaning, we will be able to draw the semantics/pragmatics distinction in a way that preserves all the desirable criteria. I will start by saying what those criteria are, then move on to showing why the standard view has been led to give up some among those. I will also show why indexical pronouns are quite essential to

---

[9] A number of collections of articles specifically on this issue may be mentioned: Turner (ed.) (1999), Bianchi (ed.) (2004), Szabo (ed.) (2006), as well as two special issues of the journal *Synthese* vol. 128, 1-2 (July 2001) and Stojanovic (ed.) (forthcoming).

the problem. Finally, we will see that the semantics/pragmatics distinction becomes again "nice and neat" once we accept that the semantic content of indexicals is their lexical meaning.

1.4.1. The Semantics/Pragmatics Distinction in the Early Days

Before indexicality and other semantically interesting phenomena of context-dependency came to bear on the discussion of the semantics/pragmatics distinction, we may speculate that there was some sort of consensus on what the distinction was, and that the following criteria were believed to converge unproblematically towards a single distinction. The criteria may be roughly formulated as follows:

(i) semantic elements are **lexically encoded** in the meaning of the words, pragmatic elements are not lexically encoded;

(ii) pragmatic elements are those that deploy various **contextual factors**; semantic elements are context-insensitive;

(iii) semantic elements have a bearing on **truth and truth conditions**; pragmatic elements do not affect truth conditions;

(iv) semantic elements obey **compositionality**; pragmatic elements are not necessarily compositional.

To get a better understanding of the motivations behind the four criteria, let us look at an example of something that is uncontroversially pragmatic. Suppose that I go into Prof. Cheng's office, and I say:

(1) It's warm in here.

I would typically utter a sentence such as (1) not necessarily to *inform* the hearer that it is warm in that place – plausibly, Prof. Cheng may tell for herself whether or not it is warm – but to *convey* something along the lines of:

(2) I wish you'd open the windows.

What I would thereby convey – which is what I would have *expressed* if I had uttered the sentence in (2) – is uncontroversially only a *pragmatic* element with respect to (1). And indeed, we can see that it turns out to be pragmatic according to the four criteria:

(i) my desire that my interlocutor opens the windows is **not** lexically encoded in the meaning of sentence uttered in (1);

(ii) in order to convey the desire in (2), I must rely on various **contextual factors**; my interlocutor must reason about what my intentions were in uttering (1), and so on;

(iii) the desire that I thereby convey has no bearing on the **truth** of my utterance of (1) – (1) is true iff it is warm in there, whether or not I actually wish that my interlocutor would open the windows;

(iv) what I convey does not enter the possible **compositional derivations** that one can perform on the sentence uttered in (1); for instance, "It's **not** warm in here" does not necessarily convey that I **do not** wish you'd open the windows.

So far so good. But things start getting complicated as we start encountering elements that according to some criteria appear to count as *semantic*, while according to other criteria, they come out *pragmatic*. And the case at point is precisely indexicals. We don't even need to bring in another example – just reconsider (1). The place of which I am talking about and for which I am saying that it is warm in there, is that semantic or pragmatic?

1.4.2. Semantic or Pragmatic? The Case of 'Here'

According to the criterion (i), taken at face value, the locational element in (1) – Cheng's office – is a *pragmatic* element. For, the lexical meaning of 'here' does not *encode* that the word should stand precisely for Cheng's office. If it did, then every time I used the pronoun 'here', I would need to be talking of Cheng's office, which is absurd.

Of course, one could tamper with the idea of "lexical encoding". For instance, for someone like François Recanati, the mere fact that there is a *word* that appears to stand for Cheng's office – the word 'here' – and that the lexical meaning of this word invites you to search for a relevant place – presumably, the place of utterance – is sufficient to make Cheng's office "lexically encoded" in the sentence in (1). Be this as it may, what remains uncontroversial is that the place itself – Cheng's office – is not *part of* the lexical meaning of 'here'.

Cheng's office does not turn out to be a semantic element of (1) according to criterion (ii) either. Obviously you need context in order to connect that place with the sentence in (1). You need to look into the context in a quite banal way – you only need to see what the context *is,* that is, what its location is – but that still counts as a contextual factor. So, according to criterion (ii), Cheng's office is only a pragmatic element in (1).[10]

But now, when you take criterion (iii) – and you take it at face value – Cheng's office would seem to have to be a *semantic* element in (1). For, the truth of (1) obviously depends on how things are with that very place. If it is not warm in there, in Cheng's office, then (1) is false, regardless of whether it is warm anywhere else. So the place is obviously relevant to the truth value of (1). *Ergo,* you might say, the referent of 'here' is a semantic element of (1). Criterion (iv) similarly suggests that the place is a semantic element, since it is preserved in the compositional derivations that (1) goes into: e.g. the negation of (1) ascribes lack of warmth to that same place, "It is possible that it should be warm in here" ascribes a possibility of warmth to that same place, etc.

The standard response to the case of indexicals is to hold onto criteria (iii) (truth-conditional impact) and (iv) (compositionality), and to give up criteria (i) and (ii), taken at face value. But since there is still something plausible about the two criteria that the standard view is being led to give up, rather than just give them up, why not try to amend them? I have already pointed

---

[10] In the next subsection, we will see how one can tamper with the notion of contextuality as well.

out how one can reinterpret the idea of lexical encoding in such a way that Cheng's office ends up being "lexically encoded" in the meaning of (1). And, as for criterion (ii), one might want to allow for "semantic contextuality" – that is, for the idea that certain semantic elements depend upon the context. Precisely, indexicals are thought to contribute *contextual* elements to the *semantic* content. But let us see more closely where this assumption leads.

1.4.3. "Good" and "Bad" Context

Indexicality has opened the doors of semantics to contextuality. However, the general feeling seems to be that one needs to place a guard on those doors. If you let too much context in, it will mess up your semantics. So the guard should only let in those contextual elements that are "invited" – invited into semantic content by some syntactic element and its lexical meaning jointly. Let us illustrate the idea by considering one of our working examples:

(3) She is dangerous.

In the syntactic structure of the sentence in (3), there is an element – the pronoun 'she' – whose lexical meaning "invites" an element from the context. The lexical meaning tells you to look into the context of utterance and search for some salient female. Suppose that I utter (3) while clearly referring to Suleïka. Then, on the standard story, Suleïka is an element that, sure, comes from the context, but is called for by the pronoun 'she' and its lexical meaning, which makes her a "good" element. And, in general, referents of pronouns and indexicals, though obviously fetched from the context, are thought to belong to the "good" part of the context, and to be therefore admissible into semantics.

Now might be a good time for a digression that would have had to come at some point or another. It concerns the difference between the so-called "pure" or "automatic" indexicals, and the rest. "Pure" indexicals are thought to be those indexicals whose lexical meaning is sufficiently rich and precise to determine, all by itself, the indexical's reference. Thus 'she' is obviously

"impure": its lexical meaning is insufficient to decide who is the person referred to with the help of the pronoun on any given occasion. For instance, if there are several salient women in the situation in which I utter (3), my interpreter will typically need to make "a best guess" in trying to decide which one among them I can plausibly be talking of. On the other hand, it is often held that the indexical 'I' is a pure indexical, since its meaning alone, on any given occasion, uniquely determines a referent (viz., the speaker). I am bringing in this distinction now, because it is sometimes thought that what makes some contextual element "good" (that is, good enough to be admitted into semantic content) is that, given a context, it gets determined in such an "automatic" way, uniquely determined by some syntactic element's lexical meaning. But if automaticity is to be a criterion of "semantic contextuality," then so much the worse for those who would want to amend criterion (iii) in this way. The idea that there are automatic indexicals appears to be pretty much hopeless. Rather than argue for this, I shall refer the reader to Predelli (2005: 40-57).[11]

Now, let us turn back to that guard at the semantics' gate. He knows what he needs to do: let in those contextual elements that have been invited by some syntactic element (and its meaning), deny entrance to anything else. But can we be sure that even then, he won't let too much context in, which would mess up our semantics? Some, like Cappelen and Lepore (2005), have argued that once you start letting context into semantics, you get onto a slippery slope: there is no way for you to stop letting more context in. They believe, though, that the slippery slope only starts right after indexicals; perhaps with quantifier domain restriction, or with gradable adjectives, or with the location parameter of weather predicates such as 'rain'. Whatever their argument

---

[11] Predelli attacks the idea that "utterances are always correctly represented by means of indexes whose co-ordinates correspond in an obvious manner to the parameters of the context of utterance" (41), such as, for example, the idea that an utterance of "I live here" must be correctly associated with an index whose co-ordinates are *the speaker* and *the place of utterance* (as would follow from the claim that 'I' and 'here' are "pure" indexicals). He presents a variety of examples, then convincingly argues that they cannot be handled correctly by the view that assumes the automaticity of 'I', 'here', 'today', etc.

might be worth, what is more interesting for us here is the idea (not theirs, to be sure) that as soon as you let *any* context in – as by allowing the reference of indexical pronouns – you might be already on a slippery slope!

1.4.4. A Slippery Slope?

It is not my goal to try to run a slippery slope argument against the standard view. All that I wish to do at this point is to help my reader see that if our guard has been taught that Suleïka is a good element and that he should let her into the semantic content of (3), then he will find it quite difficult to turn away certain other contextual elements. And if he lets those in, then he won't know how to say 'no' to others, and will end up holding the doors of semantics wide open to context and pragmatics.

Thus suppose that, pointing at a watermelon, I say:

(4) This watermelon is red.

With watermelon, one normally doesn't know what its color is until one cuts into it – it might turn out to be red, or pink, or yellow. In fact, people sometimes make bets as to the color of the watermelon. So imagine that in a situation like that, just before cutting it, I utter (4). Is my utterance of (4) true or not? It depends on what we'll see inside. Suppose that when we cut it, we see that it is indeed red. Then (4) is most obviously true.

What does that have to do with the semantics/pragmatics distinction? The following. Imagine now that Inma is painting a still life, and that she needs for her composition a biggish fruit of roundish shape that is red – red in the way in which, say, a red apple is red. Suppose that in this situation, in reference to the same watermelon as in (4), I say:

(5) This watermelon is red.

It seems that I will be now saying something false with (5). So we have the same sentence as in (4), uttered in reference to the same watermelon, the

relevant worldly circumstances are the same (the watermelon did not change color in the meantime) but one sentence is true and the other is false.

One might try to look for a way of deriving a semantic content for (4) that is going to be different from the semantic content of (5) – and if (4) and (5) have different contents, then of course they may have different truth values. One way of getting different contents would be to make the word 'red' contribute differently to the two contents. For example, in (4), the word 'red' might contribute the property of being *red from the inside*, a property that this watermelon has indeed, while in (5), in might contribute the property of being *red from the outside*, a property that the watermelon lacks.

Suppose that we do as has just been suggested. Then the way of being a certain color becomes an element of semantic content. Different ways of being red, in different contexts, reach into the semantic content. Such ways of being red are contextual elements, and since they eventually reach in the semantic content, you would think that they are "good" elements. But are they? Let's see: is there a syntactic element that "invites" the way in which the watermelon is red into the semantic content? Well, yes, presumably the word 'red' itself. But is it part of the lexical meaning of 'red' to invite you to look at the way in which the thing is red? It does not seem to be, but on the other hand, what could prevent us from thinking of the lexical meaning of 'red' in that way? After all, in the same way in which the lexical meaning of 'she' only constrains the referent to being a female, the lexical meaning of 'red' could similarly constrain its referent – a certain property – to being a certain way – say, redlike, but need not determine any unique such property. So for example, both being red from the outside and from the inside are properties of being red in a certain way, so both can be what 'red' contributes to semantic content, but they are still different properties and the contents of which they are part will therefore be different contents.

At this point, defenders of the mainstream views (including minimalist) would likely attempt to show that the way in which the pronoun 'she' invites

Suleïka into the semantic content of (3) is radically different from the way in which 'red' might "invite" a certain way of being red (from the inside or from the outside) into the contents of (4) and (5). That the two are different I have no doubt. After all, one is a personal pronoun, the other is an adjective of color. But what is less clear is that in the case of 'she' the appeal to context is semantically monitored in a way in which it is not in the case of 'red'. In both cases, the lexical meaning of the word is not sufficient to determine what the speaker meant to express; that is, which female the speaker said was dangerous, and which way of being red the speaker said the watermelon was. In both cases, the lexical meaning only constrains what the speaker could have reasonably meant to be saying. In the one case, it constrains the person said to be dangerous to being female, in the other, it constrains the color the watermelon is said to have to being red one way or another.

The problem is that both philosophers and semanticists are generally reluctant to accept the idea that a color adjective such as 'red' comes with an implicit argument that takes as its value a certain way of being red. They are also reluctant to accept the idea that the *way* in which a red thing is red reaches into the semantic content of sentences containing the word 'red', and that it does it (partly) in virtue of the lexical meaning of 'red'. So, if one wanted to distinguish between the ways in which the two expressions, 'she' vs. 'red', "appeal to context" and to draw a boundary between the two that would make the contextual contribution "good" in the former case and "bad" in the latter case, I submit that the boundary will have to rest upon some arbitrary decisions.

1.4.5. *The Semantics/Pragmatics Distinction, Nice and Neat Again*
The view on the semantics of indexicals that I am proposing allows us rehabilitate the four criteria for the semantics/pragmatics distinction in their original form, and divide the territory between semantics and pragmatics in a way as neat as it can get. Before I show how this may be done, I would like to

point out that such a nice and neat distinction may already be drawn in the view recently proposed in Predelli (2005). Predelli's primary goal is to defend the traditional truth-conditional semantic paradigm against contextualist assaults. He goes on at some length to explain how semantics works. Semantics – or the "semantic interpretation system" – takes two inputs, and returns an output. One input is a sentence.[12] The other input is what Predelli calls an "index" – and that is a series of values for a certain number of parameters, namely a world, a time, and a series of objects – pretty much something like an assignment of values to the free variables. What the system returns as the output is what Predelli calls a "t-distribution", which is a function that maps "points of evaluation" to truth values. Now, the semantic/ pragmatic distinction, as I read it into Predelli's view, gets very neat in that what is *semantic* is the function that for any sentence-index input returns a certain output, namely a t-distribution.

Unfortunately, Predelli keeps Kaplan's terminology insofar as the term 'content' is concerned, so his "contents" are just t-distributions. But what I would call the *semantic content* (associated with some sentence) would be that function which takes an index and a point of evaluation and returns a truth value, or, equivalently, that function which takes an index and returns a t-distribution. The output, that is, the t-distribution itself, is not, in my view, semantic content.

At any rate, what is nice about Predelli's picture is that the determination of all these inputs (sentence, index, and point of evaluation) can get as pragmatic as you wish. The determination of the sentence-input will typically involve lexical and syntactic disambiguations, which require context and considerable pragmatic reasoning. The determination of the index-input will similarly require considerable pragmatic reasoning. The index decides which individual

---

[12] More precisely, Predelli thinks that this input is rather a certain *representation* of a sentence, which he calls a "clause". That is because he wants to say that a sentence can still be syntactically and lexically ambiguous, but the input to the semantic system ought to be disambiguated. Since nothing crucial hinges on this, let me just continue talking as if it were sentences that are directly fed into the interpretation system.

is to be associated with which pronoun, which time is "now", which place is relevant to the truth of the sentence, and so forth. But the determination of the relevant index is not in itself a semantic matter. Neither is that a task incumbent upon semantics, nor does the choice of the relevant index have any impact upon (what I call) the semantic content. To be sure, different indices will give you different *t*-distributions, and so you may end up with different truth values. But that does not mean that the semantic content itself will vary with the index.

Technically, the formal account that I will opt for (Appendix 1) will be similar to Predelli's, and the semantics/pragmatics distinction that I have just drawn on his behalf may be transfered over to my account. But I now want to end this section by turning briefly to the initial four criteria and by showing how they may be rehabilitated.

Remember that my proposal is that lexical meaning just is semantic content, both for indexical and non-indexical elements of language. That immediately takes care of criterion (i), on which semantics deals with what is lexically encoded, that is, with lexical meaning. It also takes care of criterion (ii) – lexical meaning does not vary with the context, nor is any contextual element going to be part of the lexical meaning. Since contextual elements do not belong to the lexical meaning, context does not even reach into semantics: all contextuality is *pragmatic* contextuality. So the only potentially problematic criterion is criterion (iii):

(iii) semantic elements have a bearing on **truth and truth conditions**; pragmatic elements do not affect truth conditions.

We do not want Suleïka herself to be part of the semantic content of the following sentence (uttered in reference to her):

(6) She is dangerous.

However, the truth of (6) clearly depends on Suleïka herself, and on whether or not *she* is dangerous.

But this is where the standard view is guilty of confusion, and where my earlier remarks about truth conditions not being limited to *worldly* conditions become relevant. Yes, Suleïka is surely relevant to the *truth* of (6). If you don't give Suleïka (or someone else) as an input to the semantic content of (6), you won't get a truth value, and *ipso facto* you won't get a truth value that will depend on whether or not Suleïka dangerous. If you give Maria as an input, you might get a different truth value than what you will get if you give Suleïka. So all of this holds. But that does not necessarily make Suleïka part of the truth *conditions*. Someone – Suleïka, Maria, whoever – is ultimately needed for the truth *value*, but not for the truth *conditions*. And what criterion (iii) says is that pragmatic elements do not affect truth conditions, but that does not say that they may not be required for the truth value. There are other contextual and pragmatic elements that can be required for truth value: for instance, if the sentence contains the word 'bank', you need to first decide which meaning to associate with this word – financial institution, or river bank? Yet lexical disambiguation is obviously not a semantic matter. Or, on the other side of the spectrum, take any sentence you wish. Its truth value will depend on what the world is like. But that obviously does not mean that the world itself is part of the semantic content of the sentence. Determining the relevant possible world is required for the truth value, but not for the truth conditions. So to sum up, once we lift the confusion from the questions of whether an element affects the truth value or whether it affects truth conditions, we see that there is no need to assume that the person talked about in a case such as (6) – in our case, Suleïka – should participate in the truth conditions.

As for criterion (iv) (compositionality), the fact that the Kaplanian notion of content is compositional does not mean that this must be the *semantic* content. Besides, Kaplan's notion of character also obeys compositionality. Now, if we conceive of semantic content as of a function from (index, point of evaluation)-

pairs to truth values, the semantic content of the following sentence (uttered in reference to Suleïka) will be compositionally determined from the semantic content of (6) and the meaning of the modal operator 'it is possible that':

(7) It is possible that she is dangerous.

As a first approximation, we will have that (7) is true at (w, t, x) iff there is some w' (appropriately related to w) such that (6) is true at (w', t, x). Neither (6) nor (7) is associated a semantic content that in any way involves Suleïka. Rather, Suleïka is to be thought of as the intended value for x, that is, as the individual at whom (6) and (7) will normally be *evaluated* for their truth value.

If we say that the reference of an indexical (by which I simply mean that thing or individual referred to by the speaker who is using the indexical) is not part of the semantic content, then we can define the semantics/pragmatics distinction in a way that reconciles the four original criteria. Here, then, is a characterization of semantics that emerges from the four criteria taken together:

(definition of semantics) Semantics deals with elements that are encoded in the lexical meaning of words, that do not depend on any contextual factors, and that are truth-conditional and compositional.

1.5. Double-Indexed Semantics vs. Two Levels of Meaning

In this section, I would like to set straight a couple of things concerning the issue of what a good semantics for indexical languages should look like. The point towards which I will be driving is that it is fine, and even desirable, to use double indexing, but that from mere double indexing to a distinction between "two kinds of meanings", such as Kaplan's characters and contents, there is a big leap, and that any reasons for making such a leap must come from considerations that go beyond the task of providing the right truth conditions for sentences containing indexicals.

1.5.1. Some Preliminaries

Although this section addresses issues related to the question of what the correct semantics for *English* ought to be like, the formal nature of these issues makes it desirable that we first get clear on the question of what a semantics for a *formal* language, such as the languages of various logics, ought to be like. A semantic framework for a language (natural or formal) may be thought of as a system such that, if you give it a sentence, it will return a truth value. But of course, before it can return a truth value, it will ask you for more. First and foremost, it will ask you for a structure of interpretation. In the case of first order logic (FOL), the structure of interpretation is simply some set of individuals U, called the universe, and a function from non-logical symbols of the language to either the elements of U (for the interpretation of individual constants) or to subsets of U (for the interpretation of monadic predicates) or to sets of n-uples of individuals from U (for the interpretation of n-place predicates). Let us call this function *the interpretation functions.*

The semantic system "knows the meaning" of logical constants: FOL knows the meaning of words like 'and', 'not', 'every'; quantifier modal logic (QML) also knows the meaning of 'necessarily' and 'possibly'; tense logic knows the meaning of expressions like 'sometimes in the past', etc. On the other hand, a word like 'pencil' is not a logical constant for FOL. Its "meaning" is only given by a structure of interpretation, and the interpretation of such a word is just some subset of the structure's universe. In any structure that adequately interprets English, it will be the set of all pencils and of pencils only.

For example, a sentence like "Maria owns a pencil", formalized in FOL by $\exists x(pencil(x) \wedge owns(Maria,x))$, will be true in some structures and false in others (e.g., it will be false in any structure that interprets 'pencil' by the empty set). Now, even in FOL, the system sometimes cannot return a truth value even when you give it a structure of interpretation. That happens, of course, if the input sentence contains free variables. And that is presumably what will happen, too, if you give the semantic system for natural language a

sentence that contains (unbound) pronouns. Take the sentence "He is a popular Italian singer". Relative to one and the same structure of interpretation, this sentence may be true and false. For instance, in "our" structure it will be true relative to the *assignment function* that assigns Toto Cutugno to the free variable that stands for 'he', and false relative to the one that assigns John Perry to the same variable.

The next tiny complication comes with the treatment of modalities. Structures of interpretation of QML need to have not only a universe of individuals, but also a set of possible worlds (and an accessibility relation on that set). The interpretation function will now send an n-place predicate not merely to a set of n-uples of individuals from U, but to a *mapping* from possible worlds to such sets. So, when you give the system a sentence, and you expect a truth value in return, the system will ask you that you give it a structure of interpretation, an assignment function, *and a world of evaluation*. So take again our sentence "He is a popular Italian singer", fix the structure, and fix also the assignment function: let Toto Cutugno be the value assigned to the variable 'he'. Then the sentence will be true when evaluated at some possible worlds – say, those that are like ours – and false when evaluated at other possible worlds – say, those where Toto Cutugno is just a pizzaïolo.

Finally, when temporal expressions are treated in the tense logic-style, in the way in which we have just seen that modalities may be treated, then the structure of interpretation will include not only a universe and a set of worlds (with their accessibility relation) but also a set of times (with their own accessibility relation), and the truth of sentences will be relative to a structure, an assignment function, a world, and a time of evaluation.

As a last preliminary, let us convene on the following interchangeable notations to mean that sentence φ is true with respect to structure S, world w, time t, and assignment f of values to the free variables:

$[[\varphi]]^{S}_{w,\,t,\,f} = \mathsf{T}$ $\qquad\qquad\qquad$ $S, w, t, f \vDash \varphi$

So, for instance, the recursive truth definitions of the necessity operator $\Box$ can be expressed in either way (R being the accessibility relation):

$[[\Box\varphi]]^S_{w,\,t,\,f}= \mathsf{T}$ iff   for every w' s.t. wRw', $[[\varphi]]^S_{w,\,t,\,f}= \mathsf{T}$

S, w, t, f $\vDash \Box\varphi$ iff   for every w' s.t. wRw': S, w', t, f $\vDash \varphi$.

So much for the preliminaries.

1.5.2. *Two Motivations for Double Indexing*

Let's start with the temporal indexical 'now'. Consider the following:

(1) Maria said that she would call.

(2) Maria said that she would call now.

As Hans Kamp (1971) pointed out with a similar example, the word 'now' in English does not appear to be semantically vacuous. Clearly, (1) and (2) have different truth conditions. Now, the reason why the presence of 'now' in (2) is a problem for simple tense logics is that, at a first glance, only definitions on which 'now' is a vacuous operator seem available – and in that case, you get the same truth conditions for the two sentences, which, using semi-formal English, may be put as follows:

(3) sometime-past (Maria says (sometime-future (Maria calls)))

But (3) only captures the sense of (1). To capture the sense of (2), we need to put an indexical operator now in front of the sentence 'Maria calls', and to define this operator appropriately. What does 'now' do in (2)? It takes you back to the time at which you are evaluating (2) as a whole. The suggestion made by Kamp, adopted by Kaplan and Lewis and nowadays by almost everyone, was to make the definition of truth relative to *two* time coordinates rather than a single one. The following definitions illustrate the idea:

S, w, t_1, t_2, f $\vDash$ sometime-past φ iff   for some t<t_1: S, w, **t,** t_2, f $\vDash \varphi$

$$S, w, t_1, t_2, f \vDash \varphi \quad \text{iff} \quad S, w, \mathbf{t_2}, t_2, f \vDash \varphi$$

It should be noted that Kamp explicitly says that the indexical 'now' requires double indexing in tense logic, but that if you use variables for times, then you can easily capture the relevant sense of (2).[13] And indeed, here are two formulae of FOL that may respectively represent (1) and (2), provided that you assign to the free variable t_0 the time that is the present time of (1) and (2):

(4) $\exists t < t_0$ Maria said-at-t ($\exists t' > t$ (Maria calls-at-t')).

(5) $\exists t < t_0$ Maria said-at-t (Maria calls-at-t_0).

Let us now move to our second motivation for double indexing. Consider the following sentence:[14]

(6) Someday everything that is flourishing will be faded.

This natural language sentence is ambiguous. On one reading – the non-charitable reading as far as the sentence in (6) is concerned – it says that at some point in future, everything that is *then* flourishing will *then* be faded. Note that single-indexed tense logic has no problem accounting for that reading. On another reading, (6) says that if something is flourishing (now) then at some point in future, that thing will be faded. Simple tense logic is able to account for that reading as well (though at the cost of departing from the surface syntax of the sentence in (6)):

(7) $\forall x(\text{flourishing}(x) \rightarrow \text{sometime-future}(\text{faded}(x)))$

But arguably, there is yet a third, "collective" reading, on which (6) says that there will be some point at which everything that is flourishing now is faded – faded *together*, if you wish. This third reading entails (7), but note that

---

[13] Cf. Kamp: "I of course exclude the possibility of symbolizing the sentence by means of explicit quantification over moments. (...) Such symbolizations, however, are a considerable departure from the actual form of the original sentences which they represent – which is unsatisfactory if we want to gain insight into the semantics of English. (1971: 231f.).
[14] The motivation comes from Crossley and Humberstone (1977), who make the point on an example that involves modal logic and the indexical 'actually'.

(7) may be true even if there is no unique time at which everything now flourishing is faded. (Just imagine that whenever the begonias are flourishing, the hortensias are faded, and when the begonias start fading, the hortenisas flourish, and so on.)

The suggestion, then, is to use the indexical 'now', with its double-indexed semantics, in the representation associated with that third reading of (6):

(8) `sometime-future` $\forall \mathtt{x}(\mathtt{now}(\mathtt{flourishing}(\mathtt{x})) \rightarrow \mathtt{faded}(\mathtt{x}))$

I leave it as an exercise to the reader to coin examples analogous to the previous one that, instead of the temporal indexical 'now', use the modal indexical 'actually', and that similarly motivate the use of *two* rather than a single *world* coordinate. What matters is that, at the end of the day, we will have the following truth parametrization:

$S, w_1, t_1, w_2, t_2, f \vDash \varphi$

The first index – the pair (w_1, t_1) – is what we think of as the point of evaluation; those are the coordinates on which the usual, non-indexical intensional operators (such as 'necessarily') operate. The second index – the pair (w_2, t_2) – is what we think of as the "designated" world and time; those are the coordinates that indexical operators deploy in their truth clauses. What an indexical operator does is, roughly, to set the value of the first index to the value taken from the second index.[15]

I would like to add that in the formal framework that I am going to propose (cf. Appendix 1), the truth parametrization will be exactly as above – nothing more. In this respect, the double-indexed semantics that I am endorsing is significantly different from both Kaplan's and Lewis's, in both of which the second index is more complex: it is a "context" that, besides the time and the world, contains the *agent* coordinate. (For Kaplan, a context is a quadruple

---

[15] Technically, you could also have operators that operate on the second index, in pretty much the same way in which intensional operators operate on the first index. Such operators are called *monsters*.

agent-place-world-time, but note that the place may be *defined* as the place where the agent is located at the time and at the world of the context. For Lewis, it is a triple agent-world-time, also called a *centered* world.) I can dispense with the agent coordinate because I will be treating the indexical 'I' and other pronouns by means of free variables – but this is already anticipating too much. Finally, since I have talked about Predelli's framework in the previous section, let me mention that what he calls "index" would, in terms of the truth parametrization proposed above, correspond to the second index (w_2, t_2) together with the assignment f, while what he calls "point of evaluation" would correspond to the first index (w_1, t_1).

1.5.3. Kaplan's Motivations

In the work of Kaplan, we find neither of the two motivations that I have discussed in the previous subsection. We find somewhat similar motivations, which also have to do with embedded indexicals. However, Kaplan seldom distinguishes motivations for mere double indexing, as it were, from motivations for *two sorts of meaning* (such as characters vs. contents). The lack of a clear separation between the two sorts of questions (the more technical vs. the more conceptual) can be observed in many places in the *Demonstratives*. For example, Kaplan (1977: 499) invites his reader to consider this sentence:

(9) It is possible that in Pakistan, in five years, only those who are actually here now are envied.

He then rightly rejects the interpretation provided by single-indexed theories that construe indexicals as vacuous operators. However, he does not even consider an interpretation which *is* available to single-indexed theories, and which would go as follows (it is analogous to (7) *qua* a rendering of (6); let us, for the sake of simplification, leave the indexical 'here' unanalyzed – nothing important hinges on that):[16]

---

[16] You will also notice that the order of entailment is reversed – that is the effect of 'only' in "*only* those who are here now".

(10) $\forall x(\Diamond \text{in-five-year}(\text{is-envied-in-Pakistan}(x)) \rightarrow \text{is-here}(x))$

Note that in (10), we have eliminated the indexicals 'now' and 'actually' (and, by the same technique, we could eliminate 'here'). And (10) seems a fine rendering of (9), without requiring any double indexing.[17]

So what are the motivations for double indexing that Kaplan draws from the example in (9)? Here is what he futher writes:

> It may be objected that this only shows that indexicals always take *primary* scope. (...) This objection attempts to relegate all direct reference to the implicit use of the paradigm of the semantics of direct reference, the variable. (499)

He then proposes the following as a tentative rendering of (9) on behalf of the view that "relegates direct reference to the implicit use of the variable":

(11) $\exists w \exists p \exists t$ (w=the actual world $\wedge$ p=here $\wedge$ t=now $\wedge$

$\Diamond$ in Pakistan In five years $\forall x(x$ is envied $\rightarrow x$ is at p-t-w)

Normally, when you want to criticize a view, you try to present it so that it looks as plausible as possible – a methodological point that Kaplan definitely ignores when offering (11) as a tentative rendering of (9). For, the view that he is attacking is supposed to be one that has recourse to "the variable" – a free variable, not a bound one. Also, why would this view still use an intensional operator like $\Diamond$, when it can quantify existentially over the world variables? So, if we wanted to really consider the views that Kaplan says that he would consider, here is a more plausible rendering of (9) to consider:

(12) $\exists w \exists t$ (t=t_0+5 years$\wedge \forall x(x$ is envied in Pakistan at t w $\rightarrow x$ is at p_0 at t_0 w_0)

---

[17] More precisely, it is a "fine" rendering to the extent that it provides the correct truth conditions. But it might fail to be quite satisfactory to the extent that there is a mismatch between the surface syntax of the sentence in (9), where 'it is possible' appears to take wide scope over the rest, and the formula in (10), where the universal quantifier that we use to translate "only those" has wide scope over the rest.

And then, all you need to do is assign the present time, the actual world, and the place of (9), to the free variables t_0, w_0 and p_0 respectively.

Now, is there anything wrong with (12) *qua* a formal rendering of (9)? No, nothing. It is certainly not wrong in the sense that the system might return value False when you would expect to get True, or the other way round. Also, recall that Kamp explicitly acknowledged that double indexing was required only if you wanted to treat 'now' in the tense logic-style. But if you just use semantics with time variables, free or quantified, then you are no longer playing the same game – the arguments for double indexing no longer go through.

But let us go back to Kaplan. Here is what he writes:

> Such transformations [as (11)] [...] do not provide an alternative to
> Principle 2 [that indexicals are directly referential], since we may still ask
> of an utterance of (9) in a context c, when evaluating it with respect to an
> arbitrary circumstance, to what do the indexicals 'actual', 'here' and 'now'
> refer. The answer, as always, is: the relevant features of the context c. (499)

What is nice about this passage from Kaplan is that it makes it clear that Kaplan's concerns and motivations are not, as we might nowadays say, *properly semantic*. He is not worried, the way Kamp or Lewis or Crossley or Humberstone are, that however you try to define, say, the operator 'actually' in your modal logic, or the operator 'now' in your tense logic, you end up with cases in which the system returns a wrong truth value: the formula is, say, true in a situation in which the target natural language sentence is intuitively false. Rather, Kaplan is concerned with *metasemantic* issues. The question that he asks is: what do the indexicals 'here', 'now' and 'actually' *refer to?* And in asking this question, Kaplan already makes a leap from the issue of how to get the right truth conditions for sentences containing indexicals embedded under modal or temporal operators – an issue that admittedly calls for double indexing – to issues that are going to motivate his character/content distinction. For, the very idea of taking a sentence *as uttered*

in a given context and then evaluating it *at some arbitrary circumstance* is already part of the big picture that Kaplan is trying to put forward. Inherent to this idea is the notion of *content*, conceived of as that which the sentence expresses in the context of utterance and for which you can ask if it is true or false with respect to this or that circumstance. And Kaplan's further claim, in the passage quoted, is that what indexicals contribute to this content are "the relevant features of the context c".

1.5.4. The Logic of Double-Indexed Semantics

There is another well-known type of argument for double indexing. Recall Kaplan's central argument against the index-theory of David Lewis, Dana Scott, et al.[18] Kaplan asks us to consider the following sentence (as uttered by him on 4/23/1973, in LA):

(13) I am here now.

Kaplan famously observes that (13) is "deeply, and somehow universally true", in a way in which the following sentence is not:

(14) Kaplan is in Los Angeles on April 23, 1973.

Kaplan wants (13) to come out as a logically valid sentence in the logic of indexicals. And here is, in rough lines, his argument against the index- theory. Either it should be the case that the formal rendering of (13) is a logically valid sentence, hence true at whichever index it is evaluated – but then, in virtue of the rule of necessitation, the following is also true at every index:

(15) Necessarily I am here now.

Kaplan thinks that this is a bad consequence, because (15), he thinks, is false, hence certainly not *logically* true. Otherwise, the formal rendering of (13) provided by the index-theory is only true at some indices – but then, we have failed to capture the "deep and universal" truth of (13).

[18] See Kaplan (1977: 509). Elements of index-theory may be found in Lewis (1972) and Scott (1970).

The putative logical validity of "I am here now" is the best known example, but also one of the most controversial. The reason why Kaplan thinks that it is logically valid is that he thinks that its truth is warranted by the sheer lexical meaning (= *character*) of the sentence. But he is wrong in this respect. The truth of (13) is warranted by the lexical meaning only if one takes for granted the additional assumption that the speaker *must be located at* the place of utterance (=here) at the time of utterance (=now). That this is indeed an additional assumption is made clear by the fact that Kaplan, in his logic, only admits among the "contexts" of some model-theoretic structure S those quadruples (a, p, t, w) such that (a, p)∈Int^S('is located at', t, w). But if such a stipulation is required in order to make (13) come out logically valid – and it is required – then the presumed logical validity of this sentence is not convincing to begin with.[19] For, it is clearly not a matter of logic, but a merely *empirical* matter, to determine who is located where at what time. Yet, for Kaplan, this matter must be settled before you can even decide whether something is a structure of interpretation of the logic that he proposed.

So let us forget about "I am here now", and let us look at a much less controversial example that can be used equally well to make the same points that Kaplan is trying to make against the index-theory. Consider:

(16) If it's actually raining in Paris now, then it's raining in Paris.

Unlike (13), here we really have a sentence that sounds logically valid. So assume single indexing, and let the indices just be world-time pairs. Then (16) will presumably be true at any index at which you evaluate it. Hence, in virtue of the necessitation rule and of the temporal analog of the same rule ("if S is valid, then always S is valid, too"), we get it that the following sentence is also true at every index:

(17) Necessarily, it is always the case that if it's actually raining in Paris now, then it's raining in Paris.

---

[19] See also Predelli and Stojanovic (forthcoming).

But, Kaplan might say, (17) shouldn't be true at every index, because that sentence is false. Suppose that it is actually raining in Paris now. Then it follows from (17) that it's always raining in Paris in all possible worlds. But yesterday it wasn't raining in Paris. Ergo (17) is false.

Let me end this subsection by showing that if you use double indexing, that does not commit you to having (16) valid and (17) false. You can have double indexing with a logic in which the necessitation rule holds. Consider these two formulae:

(18) actually(p) → p

(19) □ (actually(p) → p)

Let us first see how double-indexed theory achieves the result of (18) being valid and (19) false. There are two tentative definitions of validity. Sentence φ is valid if, and only if, for every structure S and every assignment function f.....

(option 1)... and for all worlds w_1, w_2 and all times t_1, t_2, we have:

$$S, w_1, t_1, w_2, t_2, f \vDash \varphi$$

(option 2)... and for every world w and every time t, we have:

$$S, w, t, w, t, f \vDash \varphi$$

It does not take much to see that only option 2 will give you the validity of sentences such as (18). On option 1, (18) will not be valid precisely because the truth of 'actually(p)' will depend on w_2 while the truth of p alone will depend on w_1. Since there is no correlation between the two indices, the antecedent of the entailment in (18) may be true while the consequent is false. It is easy to check, on the other hand, that on option 1, the rule of necessitation holds.

Is there a matter of fact as to which of the two options is the right one? I believe not. Rather, this seems to be a matter of theoretical choice. The choice depends on what your priorities are. If your priority is, say, to preserve the rule of necessitation, then you should opt for 1. If your priority is, as for

Kaplan, to come up with a logic in which (18) is valid, then opt for 2. The point is that mere double indexing does not force you to go one way rather than the other.

1.5.5. Where Does This Leave Us?

My goal in this section has been to show that there are some good motivations for double-indexed semantics, motivations related to properly semantic issues, such as the question of how we should build our semantic system so that whenever we give it as input a sentence that is intuitively true (with respect to such-and-such parameters), the system's output (with respect to the same parameters) is value True. Now, those motivations, oddly enough, are not to be found in Kaplan's work. Or, if you wish, Kaplan's motivations involve a mixture of such semantic questions and what we might call *metasemantic* questions, whose very formulation presupposes such notions as the *reference* of a given indexical, the *content* of a given sentence (as uttered in a given context), *evaluating* that content *at a circumstance*, and so on and so forth.

So the lesson, to put it crudely, is that double-indexed semantics is one thing, and character/content semantics is yet another thing. Philosophers often tend to confuse the two, or, at any rate, not to keep them as clearly distinct as they ought to be kept. That is unfortunate, and it is even more so since the urge that the distinction be kept clear was made early on, and with great insight, by David Lewis (1980). He writes:

> I see Stalnaker and Kaplan as putting forth package deals. Offered the whole of either package – take it or leave it – I take it. But I would rather divide the issues. Part of each package is a preference, which I oppose as unwarranted and arbitrary, for variable but simple semantic values. (42)

Lewis here objects to Kaplan and Stalnaker that they have identified a certain level in their double-indexed semantics as *the* semantic level (Lewis calls it semantic *value* of the sentence; today we would call it semantic *content*), and that they have done so without good reason. The level so identified is the

level of Kaplanian content, which, in technical terms, is a function from circumstances of evaluation to truth values, and which may be associated with a sentence only once you fix the *context.* In our version of double-indexed semantics, the "context" corresponds to the designated world, the designated time, and the assignment of values to the free variables.

In the same paper, Lewis stresses that Kaplan's driving motivations are those of identifying the level of semantic content with the asserted content, or *what is said*. As we will see in 2.3.8, Lewis thought that the prospects of providing a theory of *what is said* by means of the notion of semantic content was pretty hopeless. He thought that the notion of *what is said* was too versatile to be captured by semantic means. On this issue, I am going to depart from Lewis and take Kaplan's side to the extent that the notion of semantic content that I would like to propose is also meant to stand for what is said, that is, for the content asserted by a speaker using the sentence (under normal circumstances). But where Kaplan and I part company is on the question of what that content is. I am going to argue that it is a mistake to identify it as a function from circumstances to truth values. If you want a notion of semantic content that can be worked into a good theory of what is said, then it is best to take contents stable across contexts, which comes close to Kaplan's notion of character. Of course, that will require argument, to which most of Chapter 2 is devoted. In the meantime, what we ought to keep in mind is that double indexing is fine, but from double indexing to "two kinds of meaning" there is a big leap, a leap that perhaps is not "arbitrary" but that is unwarranted and, I think, conceptually dangerous.

1.6. Embedded Indexicals and the Modal Argument

We have already encountered in the previous chapter "embedded" indexicals, that is, indexicals that occur, syntactically, in the scope of a modal, temporal, or some other sentential operator. But we have only seen what happens with

the simple temporal and modal indexicals, namely 'now' and 'actually'. What they do when they are embedded is that they go and fetch the *designated* time or world, and then set them as the new values for the time or world *of evaluation*. In this section, I would like to look at what happens with other indexicals, those for which there is no special coordinate in the parametrization of truth, such as the indexical 'I' (recall that we do not have any "agent" coordinate), and other pronouns ('you', 'she', demonstratives).

Time and again I emphasized that indexicals contribute their lexical meaning to the semantic content. Let us now see how this lexical meaning manages to escape the effects of modal and other operators in whose scope the indexical may lie. For, the material lexically encoded in indexicals is special in that it is always meant to apply, by default, in the context of utterance, rather than in some merely hypothetical context induced by the interpretation of such operators.

This section is structured as follows. First, I will lay down the problem posed by indexicals embedded under modal or other operators, problem known as "the modal argument". Then, I will show how we may resolve the problem, partly helped by the mechanism of double indexing. I will try to avoid getting too technical – the formal semantics will be properly presented in the Appendix. What is more, I will propose an *explanation* of why indexicals behave the way they behave. Hence, even if we think of double indexing as a technical trick, we have a background story to tell on the issue of why certain words (indexicals) typically go and look into the designated world-time index, while other words (non-indexicals) go and look into the world and time of evaluation, which may or may not coincide with the designated ones. Finally, I will look at a more sophisticated version of the modal argument – one that purportedly does not concern indexicals embedded under any operators, but can be run directly on indexicals in simple, unembedded sentences. I will argue that even in this version of the modal argument, it is still the interaction between modal expressions and indexicals that poses a problem and needs

explaining, and that an explanation along the lines of the one given for the simple version of the modal argument, which makes use of double indexing, is applicable to the sophisticated version as well.

1.6.1. The Modal Argument

It is widely believed that Kripke (1980) and then Kaplan (1977) have refuted the descriptive theory of proper names and of indexicals (that is, the theory that roughly says that these expressions are just definite descriptions in disguise). Let it be clear that the theory of indexicals that I am proposing is not a descriptive theory.[20] I am certainly not proposing that an indexical such as 'I' is synonymous with any definite description, such as 'the speaker'. For one thing, by being different *types* of expressions, indexicals and definite descriptions are to be expected to behave differently, both from the viewpoint of syntax as of semantics. Moreover, as I have emphasized, we should not expect that there will be some description that will capture the lexical meaning of a given indexical. It is only for convenience that we use descriptions such as 'female' in association with the pronoun 'she', or 'speaker' in association with 'I', or 'place of utterance' in association with the adverbial pronoun 'here'. But, time and again, those descriptions are not exactly the lexical meanings of the indexicals but are, at best, approximations to their lexical meanings.

This being emphasized, there is still something in my proposal that you might think "smells like descriptivism". It is the very rejection of the referentialist claim that the semantic content of an indexical is just its referent.[21] That, and the claim that the semantic content of an indexical is its context-invariant lexically encoded meaning, which, albeit not *reducible* to any specific definite description, is still some sort of general condition or constraint. So you might think that because of these features, my proposal

---

[20] In earlier versions of my proposal, it could happen that I would present it as a defense of some version of descriptivism. I would like to thank Neftali Villanueva for insisting that I should not present my view as a "descriptive" view (which it is not anyway).
[21] So far, I have not said anything yet about proper names. I will avoid touching upon this delicate topic until the very end of Chapter 3.

might be vulnerable to the problems that Kripke and Kaplan directed against the descriptive theory, problems known under the name of "the modal argument". I will start with the simple version of the modal argument, one which concerns indexicals syntactically embedded under a modal operator. By the end of this section, I will examine a somewhat more sophisticated version. Here is, in my own terms, how the simple modal argument might go:

> Suppose, as you do, that the semantic content of an indexical is some general condition lexically encoded in its meaning. Then shouldn't it be possible to replace the indexical with a definite description that makes this condition explicit, without thereby altering the truth value of the sentence. But consider the following pair of sentences:
>
> (1) She could have been a genius.
>
> (2) The most salient female could have been a genius.
>
> Suppose that I utter (1) in reference to Mirka, and suppose, for the sake of the argument, that it just could not have happened that she could be a genius. Then (1) is false. Now consider (2) as uttered in the same context. Arguably, (2) is true, because some actual female genius could have been most salient in that context (for instance, I could have been pointing at a photo of Marie Curie). But then, the two utterances differ in truth value. Therefore, given that (1) and (2) are uttered in the same context, yet differ in truth value, there must be some difference in their semantic content, and this difference must be traced to the semantic contributions of the pronoun 'she' and the definite description. Therefore, the semantic content of 'she' cannot just be its lexical meaning, which arguably boils down to the condition of being a salient female."

One place where this argument may be challenged is the assumption that the contexts of utterance of (1) and (2) are the same. For, they are not. In (1), I am referring to Mirka, I am talking of her, I am trying to tell you something about her. None of this happens in (2), on the reading on which the possibility

of pointing at Marie Curie makes it true. But let us put aside this difference in the respective contexts of (1) and (2). For, there remains something interesting and stimulating in the argument that I have outlined on behalf of my opponent, and it is her request that I explain why the lexical meaning of 'she' does not react to the modalities in the same way in which the lexical meaning of a definite description does, even when the material lexically encoded in those meanings is pretty much the same.

1.6.2. Back to Double Indexing: How It Works for Pronouns

Recall the truth parametrization on which we have settled:

$$S, w_1, t_1, w_2, t_2, f \vDash \varphi$$

Recall that S is a structure, consisting of a universe, a set of worlds and a set of times (with accessibility relations on each), and an interpretation function, responsible for interpreting the non-logical vocabulary. The first index – the pair (w_1, t_1) – is the index of evaluation, the second – the pair (w_2, t_2) – is the "designated" index, which is deployed in the truth clauses for indexicals. Finally, parameter f is an assignment of values to the free variables.

Though we have already given them earlier, let us, for clarity, recall the definition of the modal and temporal indexicals and intensional operators:

$$S, w_1, t_1, w_2, t_2, f \vDash \text{sometime-past } \varphi \quad \text{iff} \quad \text{for some } t < t_1 : S, w_1, t, w_2, t_2, f \vDash \varphi$$

$$S, w_1, t_1, w_2, t_2, f \vDash \text{now } \varphi \quad \text{iff} \quad S, w_1, t_2, w_2, t_2, f \vDash \varphi$$

$$S, w_1, t_1, w_2, t_2, f \vDash \Diamond \varphi \quad \text{iff} \quad \text{for some } w \text{ s.t. } w_1 R w : S, w, t_1, w_2, t_2, f \vDash \varphi$$

$$S, w_1, t_1, w_2, t_2, f \vDash \text{actually } \varphi \quad \text{iff} \quad S, w_2, t_1, w_2, t_2, f \vDash \varphi$$

This much is known from before. Now, how shall we incorporate personal pronouns – indexicals 'I', 'you', 'she', 'this' – into such a framework? And how is double indexing going to help us, if at all?

Our semantic analysis of indexical pronouns has two components. The one component is the use of a free variable. The other component is the lexical meaning of the indexical. Keeping in mind that this lexical meaning is not equivalent to any given description, it is still convenient to pretend that there is a description capturing the lexically encoded material. So, for the indexical 'I', we might use the predicate 'speaker', for the indexical 'she', we would use 'female', for 'this' we use 'salient', etc. Those are predicates, hence, they bring along an argument place – and that one will be occupied by a variable. Finally, such predicates, too, get evaluated at an index, and that is where double indexing helps: because they are indexicals, their index of evaluation is systematically reset to the designated index, unlike their descriptive near-equivalents:

$$S, w_1, t_1, w_2, t_2, f \vDash \varphi(I_x) \quad \text{iff} \quad S, w_1, t_1, w_2, t_2, f \vDash \varphi(x)$$

$$\text{and } S, \boxed{w_2, t_2}, w_2, t_2, f \vDash \text{speaker}(x)$$

$$S, w_1, t_1, w_2, t_2, f \vDash \varphi(\text{she}_y) \quad \text{iff} \quad S, w_1, t_1, w_2, t_2, f \vDash \varphi(y)$$

$$\text{and } S, \boxed{w_2, t_2}, w_2, t_2, f \vDash \text{female}(y)$$

$$S, w_1, t_1, w_2, t_2, f \vDash \varphi(\text{this}_z) \quad \text{iff} \quad S, w_1, t_1, w_2, t_2, f \vDash \varphi(z)$$

$$\text{and } S, \boxed{w_2, t_2}, w_2, t_2, f \vDash \text{salient}(z)$$

Of course, you will not use variable x if it already occurs in φ, whether free or bound. And there you go. What you will get in Appendix 1 is not any deeper or more interesting than the above (it is just spelled out more properly). So this is, in a nutshell, the semantics of indexical pronouns that I am proposing.

Let us now see how this technique works for embedded indexicals. Recall the pair of sentences deployed in the modal argument:

(1) She could have been a genius.

(2) The most salient female could have been a genius.

Here are two semi-formal sentences that will, for our purposes, work quite fine as formal renderings for (1) and (2) respectively:[22]

(3) $\Diamond$ genius(she$_x$)

(4) $\Diamond$ $\exists x(\text{salient}(x) \land \text{female}(x) \land \forall y(\text{salient female}(y) \rightarrow y=x) \land \text{genius}(x))$

There is no need to comment on (4). But let us show that (3) provides the right truth conditions for (1). So suppose that I utter (1) in reference to Mirka, and let t_0 and w_0 be respectively the time and the world relevant to the truth of my utterance of (1) – typically, those will be the time and the world *of* the utterance. Then we will presumably want to know the truth value of sentence (3) with respect to the following parameters:

$$[[\Diamond \text{ genius(she}_x)]]^S w_0, t_0, w_0, t_0, f(x)=Mirka = \;?$$

Here is the derivation that we get:

$$[[\Diamond \text{ genius(she}_x)]]^S w_0, t_0, w_0, t_0, f(x)=Mirka = \mathsf{T} \text{ iff}$$

for some w s.t. $w_0 R w$: $[[\text{genius(she}_x)]]^S \mathbf{w}, t_0, w_0, t_0, f(x)=Mirka = \mathsf{T}$, iff

$$[[\text{genius}(x)]]^S \mathbf{w}, t_0, w_0, t_0, f(x)=Mirka = \mathsf{T} \text{ and}$$

$$[[\text{salient female}(x)]]^S \mathbf{w_0}, t_0, w_0, t_0, f(x)=Mirka = \mathsf{T}$$

That is to say, (3) is true iff there is a possible world (accessible from ours) in which $f(x)$, that is, Mirka, is in the interpretation of 'genius', and if $f(x)$ – Mirka, again – is a salient female *in our world.* That's it. That's all.

1.6.3. Will It Work with Doxastic and Epistemic Operators?

Now that you have understood the gist of double indexing, you will know that the technique works perfectly fine for indexicals embedded under modal operators and under temporal operators. But you might think that indexicals

---

[22] Here we are endorsing a very old-fashioned analysis for the definite description, and ignoring the comparative 'most'. For a proper discussion of definite descriptions, we must wait until 3.1.

used in knowledge and belief reports might pose a problem for my view. For, suppose that, pointing at Mirka, whose looks may be deceiving, I say:

(5) Tarek thinks that she is a man.

You might think that on my account, the report in (5) would be analyzed as attributing to Tarek an inconsistent belief – that is to say, that it would roughly be analyzed along the lines of "Tarek thinks that x is a man and that x is a salient female in the actual world," with Mirka assigned to the free variable x.

But this objection will not go through, for the simple reason that, in all likelihood, the semantics for doxastic operators such as 'Tarek thinks that' that we would choose proceeds along the lines of the familiar Hintikka-style epistemic logic, where the truth definition of the doxastic operator deploys the possible world parameter (taking the notation '$w \succ_c v$' to mean that v is among c's doxastic alternatives for w):

$S, w_1, t_1, w_2, t_2, f \models \texttt{believes(Tarek, } \varphi)$ iff

$$\text{for every } w \text{ s.t. } w_1 \succ_{\text{Tarek}} w: S, w, t_1, w_2, t_2, f \models \varphi$$

I leave it to the reader to check that on this analysis of 'believes that', the technique of double indexing works fine. (Of course, belief reports are a difficult topic on their own, and there are certainly interesting issues related to indexicals as used in reporting beliefs. However, any further discussion at this stage would take us too much astray.)

1.6.4. A Pragmatic Explanation of the Semantic Behavior of Indexicals

In this subsection, I will try to provide a tentative explanation of why indexicals always want to be evaluated at the designated index, rather than at the "current" index of evaluation, that is, at the index induced by the interpretation of some intensional operator.[23] Recall that the designated index is, in the default case, that of the context in which the utterance is made and is

---

[23] To talk of *the* current index of interpretation is a bit misleading since, for example, when it comes to interpreting φ in the course of interpreting $\Box\varphi$ with respect to world w, then *every* world w' that is accessible from w counts as a "current" index of interpretation.

being evaluated for its truth. My suggestion, in a nutshell, is that the semantic patterns that force the interpretation of an indexical into the designated index are both optimal and enhanced by linguistic conventions.

Typically, the lexical meanings of indexicals involve relations to utterances. For example, 'I' bids you to look for someone who is speaking, 'here' bids you to look for some place where the word 'here' has just been used, and so on.[24] Now suppose that in course of a conversation, you need to interpret an utterance that contains 'I' or 'here'. You have the utterance right there at hand, and you already have the speaker and the place of utterance that satisfy the constraints under interpretation. In terms of efficiency, you will be better off if you evaluate the conditions lexically encoded in 'I' and 'here', which involve speakerhood and the property of being a place of utterance, right there in your present situation, where you already have a speaker, a place, and utterances of 'I' and 'here'. We might even posit some sort of implicit "rule" of efficient interpretation, such as the following:

> If the context of utterance (or the designated index) already supplies individuals and objects that stand in the required relations and satisfy the required constraints, then it is not necessary to look any further.

Of course, the rule may well be overriden, for instance, by Gricean maxims of quality, quantity, etc. But at any rate, this pragmatic story in terms of what is most efficient from the interpreter's point of view would give us an explanation, or at least the beginning of an explanation, of why indexicals require double indexing and why they always leap back to the designated world and time, when embedded.

I have suggested that it is most efficient to interpret the indexical in the context of utterance because the conditions lexically associated with the indexical involve properties such as speakerhood or salience, which are easy to scan for in the context of utterance. But recall:

---

[24] By way of a warning at a possible confusion, this is *not* a token-reflexive view (even if, to some, it may look like being one).

(6) It could have been that the most salient female were a genius.

Why is it that here, we are not scanning the designated index for the condition of being the most salient female (which would give us Mirka), but we only scan for those conditions at the current index of evaluation, which is just some hypothetically possible world? Isn't salience precisely supposed to be one of those properties that are easy to scan for right there, in the context at hand, that is, in the designated index?

Though this is a legitimate worry, it will dispel easily. If the speaker actually wanted to be talking of Mirka, and to say that *she* could have been a genius, he would have used the convenient, 3rd person pronoun 'she'. If the speaker preferred to use the cumbersome description 'the most salient female' instead, that must have been for some reason – probably because he wanted to mark the contrast with a possible assertion that would have been about Mirka. In other words, by going for the "marked" description over the "unmarked" pronoun, the speaker of (6) precisely indicates to the interpreter that she, too, should go for the "marked" interpretation, which, in the case at stake, is the one represented in (4).

The pragmatic explanation is not meant to predict, on a case to case basis, all the available readings of a given sentence, and especially, dismiss all the unavailable readings in the case of indexicals. Nevertheless, what was initially due to the optimization of the semantic interpretation process, and thus explainable in pragmatic terms, might well have become a conventional feature of the language. That is to say, even if you present me with some particular utterance such there is no obvious reason why the indexical should take its value at the designated index, rather than at the current index of evaluation, the patterns that force indexicals to apply to the designated index in analogous constructions have by now become so entrenched in our mechanisms of semantic interpretation that we interpret indexicals with respect to the designated index no matter what, and we do it as a matter of

our knowledge of language. In other words, the way in which indexicals require double indexing and exploit the designated index are things that one must know in order to master a language such as English.[25]

1.6.5. Back to the Modal Argument

Essentially, what double indexing does is to allow expressions buried underneath one or more operators (modal, temporal, epistemic) to ignore the current index of evaluation and only look at what's happening at the designated index. Now, if this is the main role of double indexing – and I believe indeed that it is – then two worries might arise. First, you might think that there is something arbitrary and stipulative about a semantics for indexicals that relies on double indexing. Second, you might worry that double indexing may only help solving the problems related to indexicals that actually lie in the scope of some intensional operator. For, if there is no operator to shift the index of evaluation, it makes little difference whether the indexical is interpreted at that index, rather than at the designated index. But, you might further worry, are there not problems related to the "modal profile" of simple sentences that contain indexicals but no operator?

The two worries that I have just outlined are, in fact, the same worries that philosophers have had concerning *wide-scope* theories – theories that would regard rigid designators, such as indexicals or proper names, simply as definite descriptions that always take wide scope.[26] One worry was that even

---

[25] Kaplan famously wrote: "Operators like 'In some contexts it is true that', which attempt to meddle with character, I call *monsters*. I claim that none can be expressed in English (without sneaking in a quotation device)" (1977: 511), and he also wrote that monsters "*could not be added to [English]*" (510, his emphasis). In my view, Kaplan fails to explain why there could be no monsters in English. However, if we take it to be a part of the *lexical meaning* of an indexical that it requires double indexing, and that it requires its index of evaluation to be reset to the designated index, then we might have the beginning of an explanation of why any English expression that may look like a monster would turn out to involve some sort of "quotation device".

[26] The notion of "rigid designator" comes, of course, from Kripke (1980). Not all rigid designators are "directly referential" - e.g. the definite description "the even prime number" rigidly denotes number 2, but is arguably not a directly referential expression. For a discussion of the difference between rigid designation and direct reference, see Kaplan (1977: 492-3), or Recanati (1993: 10-14).

if the trick could work from a merely technical point of view, construing indexicals or names as wide scope-taking devices was an arbitrary stipulation (cf. Kripke (1980)). And the other worry was that the trick could only work for indexicals that are embedded under some operator, but that all of the crucial points about direct reference or rigid designation could be made with respect to indexicals that come in simple, unembedded sentences.

Insofar as concerns the first worry, the previous subsection was precisely meant to show that double indexing was not just a technical trick, but that the sematic behavior of indexicals could eventually be *explained* in terms of the optimal efficiency of possible interpretation patterns. What I would like to do in this subsection is address the second worry. I will try to show that even the more sophisticated version of the modal argument, which has been claimed not to concern embedded indexicals, in fact does so. It just lets the modal operators slide in through the back door, as we might say. But let me first outline the sophisticated modal argument on behalf of my opponent:

> Suppose again, as you do, that the semantic content of an indexical is some general condition lexically encoded in its meaning. Then it shouldn't make much difference whether one uses an indexical or a definite description that simply makes this condition explicit. We had previously seen that when the expression occurs in the scope of some operator, then there is a difference, and you have shown us how the difference may be accounted for using double indexing. However, your account will only help for indexicals lying in the scope of some operator. But consider this pair of simple, unembedded sentences:
>
> (7) She is a genius.
>
> (8) The most salient female is a genius.
>
> Suppose that I utter (7) in reference to Mirka, and that I utter (8) in the same context. Suppose furthermore that Mirka isn't a genius, and could have never been one. Now, we, direct reference theorists, claim that (7)

and (8) do not have the same truth conditions. To see this, consider a possible reaction to either of those:

(9) That's not the case, though it might have been!

If (9) is uttered in response to (7), then it is false, because, by assumption, Mirka couldn't have been a genius, but if (9) is uttered in response to (8), then it is true, because some actual female genius could have been most salient in that context (for instance, I could have been pointing at a photo of Marie Curie). What this show is that (7) and (8) must differ in truth conditions, hence in semantic content, hence the semantic content of 'she' cannot just be its lexical meaning, because that arguably just boils down to the condition of being a salient female.

Just as in the simple case, one assumption on which the argument rests and that may be challenged is that of (7) and (8) being uttered in the same context. For, again, the context of (7) is one in which the speaker is talking about Mirka and referring (*qua* speaker) to her, which does not hold of the context of (8). But again, let's put this aside. First of all, note that (7) and (8) do not differ in truth value. Nor will the sentences in (7) and (8) ever actually differ in truth value, if they are uttered in the same context. But that is not what the referentialist claims – he or she claims that (7) and (8) differ in truth *conditions.* And their evidence for this claim is that the sentence in (9) will differ in truth *value* depending on whether it is uttered as a comment on (7), in which an indexical is used, or as a comment on (8), in which a definite description is used instead.

So we can already see that the problem arises only if we consider the simple sentences such as (7) and (8) within a broader linguistic context. The simple sentences themselves do not diverge in truth value, so there is no direct semantic evidence that there is any difference in their semantic content. That on which the two sentences differ is on how they interact with other *sentences*, such as (9). Now, those sentences that pose problems crucially involve device

of *sentential anaphora*, namely the pronoun 'that' in "that's not true" and 'it' in "it could have been".[27] Those among the direct reference theorists who are particularly fond of the sophisticated version of the modal argument (e.g. Soames) hardly ever consider the crucial role that such anaphoric devices play in their argument.

Now, sentential anaphora is a phenomenon still under study, and there is no general agreement on how it works. One promising approach is to always try to appropriately restore the antecedent sentence through the anaphoric pronoun.[28] In other words, depending on whether it is uttered in response to (7) or to (8), the sentence in (9) is going to be associated with different syntactic representations:

(10) It is not the case that she is a genius, though it might have been the case that she were a genius.

(11) It is not the case that the most salient female is a genius, though it might have been the case that the most salient female were a genius.

If the right approach to sentential anaphora, and perhaps to ellipsis in general, is such a syntactic one, then the sophisticated version of the modal argument collapses into the simple version and may be directly handled by such devices as double indexing (and probably by any wide scope-forcing device).

That there might be something crucially syntactic to this problem is partly supported by the fact that definite descriptions exhibit a relevantly similar behavior in constructions that involve sentential anaphora. Consider the following variation on the previous example:

---

[27] This is the same as the phenomenon sometimes called *propositional* anaphora. I have chosen the term "sentential" rather than "propositional" in order to forestall any confusion, since one might have thought that a phenomenon called "propositional anaphora" would crucially involve *propositions,* while I have stressed how propositions play no role whatsoever in my account. The choice of the term "sentential" in replacement of "propositional" is in line with the more and more frequent talk of *sentential* logic, *sentential* operators, etc. for what used to be called *propositional* logic, *propositional* operators, etc.
[28] Cf. Lappin (1996).

(12) Tarek's wife, whoever she might be, is a genius.

(13) That's not the case, though it might have been.

The right hand sentence in (13), "it might have been", is ambiguous between the reading on which it is true iff Tarek could have had married someone who were actually a genius, and the reading on which it is true iff Tarek's actual wife is not a genius but might have been one.[29] But this kind of ambiguity is very much like the regular scope ambiguity that you may get with definite descriptions embedded under a modal operator:

(14) It might have been the case that Tarek's wife were a genius.

I don't want to pursue the present discussion much longer, since in Chapter 2, we will look again, at length and in depth, at simple sentences containing indexicals and their interaction with another type of sentences that crucially involve sentential anaphora, such as "That's what she said." So, if you think that there might be any loose ends left at this stage, by the end of Chapter 2, all such loose ends will be tightened once and for all. Insofar as definite descriptions are concerned, matters will be set straight in Chapter 3.

1.7. Double Indexing vs. Presupposition

In this section, we will discuss some similarities between indexicals (and their semantics) and the familiar phenomenon of presupposition. The conclusion, however, will be that if the conditions lexically encoded in the meaning of indexicals are presuppositions, then they are a special kind of presuppositions.

---

[29] My informants did not have problems to get the two readings. Still, if one wanted to enhance the second reading, one could be explicit, for instance, by adding, "Tarek's wife has some talent, but she never had discipline and was never able to make use of her talent. If she had been raised differently, she would have been a genius." *En passant*, note that the definite description in (12) is used attributively, not referentially. This means that on the Kaplanian account, the definite description does not contribute its reference to the content of (12). It is not clear how Kaplan could account for the ambiguity of (13) and the availability of the second reading.

Unfortunately, this means that the existing theories of presupposition are not directly applicable to our purposes. Another lesson of this section is that double indexing and presupposition are mechanisms compatible with one another, but neither can replace the other.

1.7.1. Introductory Remarks

In 1.5, we saw two very good motivations for double indexing. The first was the difference in the truth conditions of the sentences "Mirka said that she'd call" and "Mirka said that she'd call *now*." Recall, however, that this was a *conditional* motivation for double indexing: it only required double indexing if there were indexing to begin with, that is, if we wanted a semantics of tense and temporal expressions along the lines of Priorian tense logic. But another perfectly acceptable reaction to this sort of cases would be to forget about index-theory and tense logic, and to use predicate logic with its good old variables and quantifiers.[30]

The other motivation for double indexing that we saw in 1.5. was the "collective" reading of the sentence "It is possible that everything that is flourishing should be faded," where we understand that there is a possible world in which everything that is actually flourishing is faded (as opposed to the weaker reading, on which for everything that is actually flourishing, there is a possible world in which that thing is faded.) Now this, too, is a *conditional* motivation in same respect in which the previous one was. It is easy to see that if you use variables for worlds, then, of course, you can express such things as there being *a single world* within which everything that flourishes *in our world* is faded.[31]

---

[30] That predicate logic (with generalized quantifiers) is more suitable to formalizing the tense system of natural language than tense logic is argued e.g. in King (200?).

[31] In fact, yet a third motivation may be added to the stock, one which does not even involve any indexicals. Consider the following:
(1) Most hostages safely arrived to their home countries.
(2) Most hostages were treated correctly.
(3) Most hostages were staying in the touristic parts of town.
The problem raised by (1)-(3) lies in the implicit time associated with the predicate 'hostage'. Only in (2) does this time coincide with the time of the reported event (they were hostages and were treated correctly at the same time). In (1), the time implicit in 'hostage' is understood as

So let us agree that double indexing is good and useful. As we saw in the previous section, double indexing may be used in solving the problem of embedded indexicals – and not only modal or temporal indexicals, but also personal and demonstrative pronouns. However, there is another mechanism that, at a first glance, appears applicable to embedded indexicals. It is the mechanism of *presupposition*.

Presupposition comprises a vast array of phenomena, phenomena that are interconnected but that still exhibit important differences. Here are just a few typical cases (with sign ⇒ used in the sense of 'presupposes'):

(1) Tarek stopped drinking. ⇒ Tarek had been drinking.

(2) It is Tarek who hit Maria. ⇒ Someone hit Maria.

(3) Tarek's wife is Cambodian. ⇒ Tarek has a wife.

The hallmark of presupposition is that not only the truth, but also the *falsity* of a given sentence, requires that the presupposition associated with some element in the sentence should obtain. If the presupposition fails, the original sentence is normally perceived as neither true not false.

In this respect, there is already something presupposition-like about indexicals. Consider:

(4) He is Cambodian.

Suppose that Tarek utters (4) while pointing at Mirka, who is a woman, and suppose that she is actually Slovakian. We might be reluctant to say that (4) is actually *false* in this case, rather than being neither true nor false. Or, suppose that Mirka actually is Cambodian. We will also be reluctant to say that (4) is true in that case. One possible explanation is that this is because of a

---

being anterior to the time of the reported event (they were no longer hostages when they arrived home), while in (3), it is understood as being posterior to the time of the reported event (they were staying in the touristic parts of town before they were turned into hostages). But guess what. This motivation, too, is conditional: it does not apply if you use (free or quantified) variables for times.

presupposition failure. The presupposition associated with 'he' is that the person referred to is *male*, which is false in the situation in which Tarek utters (4) while pointing at Mirka.

1.7.2. Projection and Cancelation

Another well-known hallmark of presupposition is their tendency to "project". Let me explain. Consider the same sentences as before, but embedded under the possibility operator. Even though the element that triggers the presupposition lies in the scope of the operator, the associated presupposition remains as before:

(5) It is possible Tarek stopped drinking. ⇒ Tarek had been drinking.

(6) It is possible that it's Tarek who hit Maria. ⇒ Someone hit Maria.

(7) It is possible that Tarek's wife is Cambodian. ⇒ Tarek has a wife.

So presuppositions, like the lexical meanings of indexicals, only seem to care about what is happening at the designated index. The current index of evaluation, such as the hypothetically possible world induced by the modal operator, is irrelevant to evaluating the presupposed material, even when the word that triggers the presupposition lies in the scope of that operator. Like indexicals, presuppositions ask to be evaluated at the designated index.[32]

Finally, presuppositions can also be "cancelled" – or, to use the modern jargon, they can be "locally accommodated". Consider:

(8) Tarek stopped drinking, if he ever drank to begin with.

(9) It is Tarek who hit Maria, if anyone did.

(10) Tarek's wife is Cambodian, assuming that he has a wife.

---

[32] Projection is not reducible to the distinction between the designated index vs. the index of evaluation. For, presuppositions also "project" when they are used with the negation operator: "Tarek didn't stop drinking" presupposes that Tarek had been drinking. The same goes for the material lexically encoded in indexical pronouns: "He isn't Cambodian" will not be rendered true by referring to someone who fails to satisfy the condition of being male, encoded in 'he'.

Is cancelability also characteristic of the lexical meaning of indexicals? That is less obvious. Nevertheless, suppose that Tarek says the following, while pointing at Mirka:

(11) He is Cambodian, if that's a man to begin with.

It seems possible to "cancel" the condition lexically associated with 'he', namely, that the person referred to should be male.

There is a difference, though, between cancelling the presuppositions as in (8)-(10), and cancelling the material lexically encoded in the indexicals. For the former, when the presupposition does not obtain, the sentence is true *by default of the antecedent;* for instance, if no one hit Maria, then (9) is ipso facto true. But consider (11). The fact that Mirka isn't a man doesn't ipso facto make (11) true. For (11) to be true, it would still need to be the case that the person referred to – in our case, Mirka – is Cambodgian.

This difference brings out an interesting feature of the presuppositions associated with pronouns (to the extent that there are such). Namely, with the usual presupposition-triggers, the presuppositions are of *propositional* character – they obtain or not relative to a world and a time, but they do not depend on any other parameters. But with pronouns, it appears that their propositions are *predicative*: to determine whether the presupposition associated with an indexical pronoun obtains, you need to ask: *of whom?* In other words, while the presupposition triggered by 'stopped' in (1) and (5) is *that Tarek had been drinking*, the presupposition triggered by 'he' in (4) is simply *being male*, and this is being presupposed *of* the person who is being referred to, namely, Mirka.

1.7.3. Concluding Remarks

Presupposition is a vast domain, where a number of phenomena that share certain feature, but still exhibit some important differences, have been grouped together. It is quite possible that indexicals are just another garden variety of presupposition-triggers. But if they are, then they are significantly

different from the rest. The presuppositions associated with indexicals are not propositional, but predicative. In other words, in order to check whether a given presupposition is satisfied, you need to have something first so that you can check whether *it* satisfies the presupposition. For example, recall the situation in which Suleïka and Tarek come to the party together, and I say:

(12) She is dangerous.

Let us assume that the presupposition associated with 'she' is *being female.* To check whether this presupposition obtains, you need to already have a potential referent – someone of whom you can ask if he or she is a female. In the case of (12), there are two potential referents: Tarek and Suleïka. And precisely, the fact that Tarek fails the presupposition – for he is not a "she" – might play a role in your deciding that I must be talking of Suleïka and that I must be referring to her in (12).

There is more than one lesson to be drawn from this section. Regarding the relationship between presupposition and double-indexed semantics, it is that presupposition cannot dispense us of double indexing. Nor will it work the other way round – double indexing by itself is not a theory of presupposition (even if double indexing might turn out to be required for a theory of presupposition). Regarding the issue of what role presupposition plays in the semantics of indexicals, we could see that if the material lexically associated with indexicals is to be treated in terms of presupposition, then it is a special kind of presupposition: one that is not propositional but predicative, and that does not just obtain or not, but obtains of some things and fails to obtain of other things.

This last point will remind you of the idea that the mechanism of direct reference does not result from the semantics of "directly referential" expressions, but, in some sense, comes first (cf. 1.2). In a sentence such as (12), I need to be able to refer to Suleïka, and you need to be able to believe that I might be referring to Suleïka, independently of the use of any pronoun or any

other linguistic device. The pronoun will constrain the range of individuals to which I can reasonably refer: the condition of being female, lexically encoded in 'she', will help you decide that in (12), I must be talking of Suleïka.

Given that there is no uncontroversial theory of presupposition that one could directly apply to indexicals, and that the available theories would have to be considerably modified for the reasons pointed out, I am going to avoid talking of presupposition for the rest of this work.

1.8. Towards a Unified Account of Deictic and Anaphoric Uses

Many among our examples involving indexicals used the third person pronouns 'he' or 'she'. Indeed, when such a pronoun is used by a speaker referring to someone, of whom she wants to communicate something to her audience, the pronoun looks like a directly referential expression *par excellence*. Someone is being directly referred to, and the conditions lexically encoded in the pronoun are either idle or help the hearer figure out of whom the speaker is talking. But it is well known that such *deictic* uses are not the only uses of pronouns. In particular, pronouns may be used *anaphorically*, and they may also be *bound*. Defenders of the mainstream view tend to simply ignore those other uses of pronouns, or, when they don't, they say things that to the linguist's ear sound very implausible. Thus Kaplan is known for having written:

> Pronouns are lexically ambiguous, having both a demonstrative and an anaphoric use. (1989: 572)

But there is no obvious reason for postulating *lexical* ambiguity. For one thing, many languages have third person pronouns that allow for the range of uses for which they allow in English, namely, deictic, anaphoric and bound

uses.[33] It would be surprising if the third person pronoun were lexically ambiguous in all these languages. Lexical ambiguity is accidental, while the deictic/anaphoric/bound ambiguity follows a pattern of regularity. So, in the best case, we would have *polysemy*, rather than genuine ambiguity.

At any rate, I believe that there is something unsatisfactory in ignoring those other uses of pronouns and simply focusing on their deictic uses. What is more, there are many expressions that are like pronouns in this respect. For instance, definite descriptions allow for deictic and anaphoric uses (though perhaps not for bound uses, but on the other hand, they have uses that pronouns arguably do not have, like the *attributive* use; cf. 3.1). Quantifiers, insofar as their domain restriction is concerned, also appear to have a range of uses that are very much like deictic, anaphoric, and bound uses. A variety of adjectives and nouns, such as 'local' or 'enemy', may also be used deictically (indexically) or anaphorically (cf. Partee (1989)). And so on.

In this chapter, I want to do what I think anyone interested in indexicals should do at some point: I want to look at the range of those uses, and at their relationship. Unfortunately, this cannot be done in a thorough and entirely satisfactory way, for reasons of space. Anaphoric pronouns are a a vast domain of study, and I will touch upon it only because it is an astonishing fact that the same words that, when used deictically, are the paradigm of indexicality, also have those other uses. I will try to show that my proposal from 1.6, provided for deictically used pronouns, also applies to anaphorically used pronouns. What is more, the semantic content of a 3rd person pronoun *remains the same* whichever way the pronoun is used. However, I must stress that if we wanted to have a unified account of deictic and anaphoric uses that works to perfection, the account on which we have settled might not be quite up to the demand. That is because we are following the truth-conditional,

---

[33] Actually, insofar as bound uses are concerned, in English and in most languages that have the reflexive pronoun, it is normally the reflexive pronoun that is used. But English does not have reflexive *possessive* pronouns, so it uses personal possessive pronouns instead; which person is used depends, of course, on the binder.

model-theoretic tradition. If we wanted a good theory of anaphora, we would probably be better off with *dynamic* semantics.[34] The reason why I have stayed within traditional semantics is that the aspects peculiar to the dynamic approach would likely obscure the main issues that I am interested in, issues that have to do with indexicals, their semantics, and their contribution to the asserted content. By staying as close as possible to the mainstream view in other respects (such as the choice of "static" over dynamic semantics), it becomes easier to see how exactly I depart from it. At any rate, those familiar with (and convinced by) dynamic semantics will know, when it comes to the formal framework, how to "dynamify" it, and nothing in my proposal should pose any new problem – quite to the contrary.

1.8.1. Deictic, Anaphoric, and Bound Uses: Some Preliminaries

It will be best to illustrate the difference among the three uses with some examples. Consider:

(1) She is dangerous. (uttered while pointing at Suleïka)

(2) Suleïka hit Maria because she had insulted her.

(3) Every woman cares for her children.

The use of 'she' in (1) is a *deictic* use, the use of 'she' and 'her' in (2) is an *anaphoric* use, and the use of 'her' in (3) is a bound use. The difference between anaphoric and bound uses is essentially a syntactic difference, and I will not go into its details. Suffice it to point out that in (2), neither 'she' nor 'her' lie, from the syntactic point of view, in the scope of their antecedents, that is, 'Suleïka' and 'Maria'. Also, the interpretation, or, as the jargon likes it, the *resolution* of anaphoric pronouns is not syntactically constrained. In (2), for instance, we understand that 'she' is anaphoric on 'Maria' while 'her' is

---

[34] A novice might think that there is such a thing as *the* dynamics semantics. The reality is that there are many dynamic semantics. Among the best known and the most influential are the *discourse representation theory* (DRT) of Kamp (1981), Kamp and Reyle (1993), the *file-change semantics* of Heim (1982), and the *dynamic predicate logic* of Gronendijk and Stokhof (1991).

anaphoric on 'Suleïka'. But this resolution is entirely context-driven. The reader will easily coin examples that, from the point of view of syntax, are isomorphic to (2), yet in which we will resolve the pronouns the other way round: 'she' as anaphoric on 'Suleïka' and 'her' on 'Maria'.

The decision whether a pronoun, on a given use, is anaphoric or bound, may not be always immediate. Consider:

(4) Suleïka is upset, because her children didn't call her.

(5) Suleïka is upset that her children didn't call her.

(6) Suleïka cares for her children.

While 'her' in (6) is bound (in fact, (6) may be view as an instantiation of (3)), in (4), both uses of 'her' ('her' in 'her children', then 'her' on its own) are anaphoric, and in (5), I honestly don't know. Yet, looking at it with a layman's eye, you might think that 'her' in 'her children' would be analyzed in the same way in the three cases.

Note also that if the antecedent is a quantifier phrase, as previously in (3), that does not guarantee that the pronoun will be bound rather than anaphoric. Thus the so-called donkey anaphora often involves pronouns anaphoric on, but not bound by, some quantifier phrase.[35] For example:

(7) If a woman comes into a restaurant and misbehaves, she will be asked to leave.

The use of 'she' in (7) is anaphoric, but its antecedent is a quantifier phrase (that is, the indefinite description 'a woman').

Finally, it has been pointed out that not only 3rd person pronouns, but also 1st and 2nd person pronouns, may have bound uses (cf. Heim (ms.)):

(8) Only I did my homework.

---

[35] On donkey anaphora see e.g. Evans (1980), Heim (1990), or Groenendijk and Stokhof (1991). The phenomenon was, of course, discovered by Peter Geach: Geach (1962).

In (8), on its most natural reading, 'my' is a bound pronoun. Only thus may we get the interpretation on which (8) means:

(9) I did my homework, and no one else did his or her homework.

So much for the preliminaries. I do not pretend to have given anything that comes close to a thorough survey of the phenomena under study. But what we have so far will do for our purposes. I should also note that in what follows, I will focus on anaphoric rather than bound uses. In particular, I will leave aside the bound uses of the 1st and the 2nd person pronoun.[36]

1.8.2. Deixis or Anaphora, Same Semantics

Let us recall the truth clause proposed for the pronoun 'she', repeated below as given in 1.6.2:

$$S, w_1, t_1, w_2, t_2, f \vDash \varphi(she_x) \quad iff \quad S, w_1, t_1, w_2, t_2, f \vDash \varphi(x)$$
$$and\ S, w_2, t_2, w_2, t_2, f \vDash female(x)$$

Whether the pronoun is used deictically or anaphorically, the truth clause remains the same. It is not difficult to see why this should so. The semantics of the pronoun, as previously said, breaks into two components: one is the use of a variable, the other is the lexically encoded condition, whose interpretation sets the index of evaluation back to the designated index. Whether the pronoun is used deictically, anaphorically, or bound, is handled through the parameter of assignment of values to the free variables. We know that this parameter plays two roles. First, it assigns values to the *free* variables – that is what happens in the deictic uses. Secondly, assignments of values to variables are deployed in the definition of quantifiers. Remember:

---

[36] A discussion of bound uses, including the bound uses of the 1st and the 2nd person pronoun, would get us entangled into issues that heavily involve syntax. Furthermore, recall that languages that, unlike English, always use the reflexive pronoun for bound uses, would not allow for Heim's problematic sentences such as "Only I did *my* homework", which, in such a language, get expressed as "Only I did *-reflexive-possessive-* homework". It is also partly because of this reason that I focus more on the anaphoric than on the bound uses of indexicals.

$$S, w_1, t_1, w_2, t_2, f \vDash \forall x\varphi(x) \quad \text{iff} \quad \text{for every assignment } f' \text{ that is like } f \text{ except}$$

at most for the value that it assigns to x,

$$S, w_1, t_1, w_2, t_2, f' \vDash \varphi(x)$$

So we can already see that what is going to distinguish the anaphoric and the bound uses of the pronoun from its deictic uses is that in the latter case, the assignment of values to the variables will be supplied directly, very much in the way in which the designated index is supplied, while in the former case, the assignment will be relevantly inherited from the interpretation of some antecedent expression.

It might be helpful to see how things work in such an anaphoric use. Recall the example from (7), somewhat simplified:

(10) If a woman misbehaves, she will be asked to leave.

Closing our eyes over the problems raised by the fact that the indefinite description occurring in the antecedent of the conditional turns itself into a universal quantifier,[37] here is, roughly, a possible derivation for (10):

$$S, w_1, t_1, w_2, t_2, f \vDash \forall x((\text{woman}(x) \land \text{misbehaves}(x)) \rightarrow \text{asked-to-leave}(\text{she}_x))$$

iff

for every assignment f' that is like f except at most for x,

if $S, w_1, t_1, w_2, t_2, f' \vDash \text{woman}(x)$ and $S, w_1, t_1, w_2, t_2, f' \vDash \text{misbehaves}(x)$,

then $S, w_1, t_1, w_2, t_2, f' \vDash \text{asked-to-leave}(x)$ and $S, w_2, t_2, w_2, t_2, f' \vDash \text{female}(x)$

Note that in the very last sub-clause, the index of evaluation is set to w_2, t_2, which is the designated index. Whether this index-resetting could do us any damage is the question to which we will turn in the next subsection. In the meantime, let us note that the truth derivation proposed for (10) gives us indeed the right truth conditions. Note that the lexically endoded condition of

---

[37] These are some of the problems that arguably require dynamic semantics. See e.g. Kamp and Reyle (1993) or Gronendijk and Stokhof (1991) for accounts in which the indefinite description is treated as an existential quantifier, wherever it occurs, yet the truth conditions of the conditional are such that this existential quantifier gets the force of a universal quantifier.

being female is idle in (10), as we might say. But that needn't be so. Even with the anaphoric uses, the lexical meaning of pronouns may play a useful role in communication. Consider the following:

(11) If a man hits a woman, she'll harm him.

(12) If a man hits a woman, he'll harm her.

The conditions of being female or male, respectively encoded in 'she' and 'he', are put at work in the disambiguation of (11) and (12). Let us, for the sake of clarity, forget about the time and world indices. Then here is what we are going to get for (11) at an initial stage:

$S, f \vDash \forall x \forall y((man(x) \wedge woman(y) \wedge hits(x,y)) \rightarrow harms(she_z, he_v))$

iff

for every assignment f' that is like f except at most for x and y,

if $S, f' \vDash man(x)$ and $S, f' \vDash woman(y)$ and $S, f' \vDash hits(x,y)$,

then $S, f' \vDash harms(z,v)$ and $S, f' \vDash female(z)$ and $S, f' \vDash male(v)$

And similarly, for (12) we will initially get:

$S, f \vDash \forall x \forall y((man(x) \wedge woman(y) \wedge hits(x,y)) \rightarrow harms(he_z, she_v))$

iff

for every assignment f' that is like f except at most for x and y,

if $S, f' \vDash man(x)$ and $S, f' \vDash woman(y)$ and $S, f' \vDash hits(x,y)$,

then $S, f' \vDash harms(z,v)$ and $S, f' \vDash male(z)$ and $S, f' \vDash female(v)$

At the next and final stage, one proceeds to *unification*:[38] the variables z and v, which are free but shouldn't be, get replaced by variables that are bound. This is where the lexical meanings of pronouns come into play. The reason why in

---

[38] On unification, see e.g. Blackburn and Boss (2003). The idea of unification originally comes from the programming language PROLOG, and, as it applies to natural language pronouns, it is mainly discussed in the field of computational linguistics.

(11) you will unify z with y and v with x, and why in (12) you will unify z with x and v with y, is that otherwise you would get something that is trivially false.[39]

Let me end this subsubsection by acknowledging that many details of the analysis have been left loose. But it was not my goal to provide, in a span of a couple of pages, a full-fledged account of anaphoric pronouns. Had I wished to do so, I would have had to look much more carefully at the syntax/ semantics interface, while I have tried to keep at distance from it. My goal has only been to point out that my main proposal, which identifies semantic content with lexical meaning, is not threatened by there being indexical pronouns that also have deictic and anaphoric uses. To the contrary, the fact that the semantic contribution of pronouns remains the same regardless of how the pronoun is used scores a point for my proposal.

1.8.3. *The Role of Double Indexing for Anaphoric Uses*

The truth clause proposed for the pronoun 'she', make us of double indexing. Now, we might want to check that this is how it should be, and that even with the anaphoric use, what is relevant is that the lexically encoded conditions ('female' for 'she', 'male' for 'he') should obtain (of the value assigned to the variable) at the designated index, rather than at the current index of evaluation. To check this, we'd better come up with a case where the current index of evaluation is different from the designated index, and where the condition obtains at only one of the two indices. Consider the following:

(13) Though Tarek is a man, it is possible that **he** should be a woman.

(14) # Though Tarek is a man, it is possible that **she** should be a woman.

The sign '#' indicates that (14) is bad. But this precisely confirms the truth clause proposed, which takes the lexically encoded condition back to the

---

[39] A tentative unification that would go as follows can only be true by default of the antecedent, which is certainly not what (11) means:

$\quad$ if S, f' $\vDash$ man(x) and S, f' $\vDash$ woman(y) and S, f' $\vDash$ hits(x,y),

$\quad$ then S, f' $\vDash$ harms(x,y) and S, f' $\vDash$ female(x) and S, f' $\vDash$ male(y)

designated index. For, if this condition were tested in the current index of evaluation, which means, in the possible world in which Tarek, having changed sex, is a woman, the condition would fail to obtain, hence (13) would be false. And, on the other hand, the condition encoded in 'she' obtains of Tarek in the possible world in which Tarek is a woman, hence (14) should have been fine if this condition were tested in the current index of evaluation. The data therefore strongly suggest that indexicals, even when used anaphorically, require that the condition that they lexically encode be tested back at the designated index.

Here is just another example that confirms the same point:

(15) If Tarek were a woman, Aaron would marry him.

(16) # If Tarek were a woman, Aaron would marry her.

Again, in the possible world in which the antecedent holds, Tarek satisfies the condition lexically associated with the pronoun 'she'. The fact that it is (15) rather than (16) that one ought to say shows that it is the designated index, rather than the current index of evaluation, that matters.

Finally, examples that involve the time coordinate, rather than the world coordinate, again confirm the double indexing hypothesis:

(17) After Tarek's surgery, everyone will like him.

(18) # After Tarek's surgery, everyone will like her.

At the future time at which Tarek will have changed sex, he will satisfy the condition of being female. Yet (18) is bad. This suggests that the lexically encoded condition is evaluated at the designated time, rather than being evaluated at the current time of evaluation.

Unfortunately, things are not that simple. Recall (10), but now, consider the sentence as embedded under a modal operator:

(19) Necessarily, if a woman misbehaves, she will be asked to leave.

The derivation that we are going to get is as follows:

$$S, w_1, t_1, w_2, t_2, f \vDash \Box \forall x((woman(x) \wedge misbehaves(x)) \rightarrow asked\text{-}to\text{-}leave(she_x))$$

iff

for every assignment f′ like f except at most for x, and for every w s.t. $w_1 R w$,

if $S, w, t_1, w_2, t_2, f′ \vDash woman(x)$ and $S, w, t_1, w_2, t_2, f′ \vDash misbehaves(x)$,

then $S, w, t_1, w_2, t_2, f′ \vDash asked\text{-}to\text{-}leave(x)$ and $S, w_2, t_2, w_2, t_2, f′ \vDash female(x)$

The problem with this derivation are those individuals who are women at some possible worlds, but not at the designated world. For, suppose that Tarek had changed sex in world w_5. Intuitively, (19) implies that in w_5, if Tarek (who is a woman there) misbehaves, he will be asked to leave. But that is not what the derivation above gives us.[40]

What this suggests is that on some of their anaphoric uses, the conditions lexically encoded in the pronouns will be evaluated at the current index of evaluation, rather than at the designated index. So the derivation that we will get, then, for (19), is as follows:

$$S, w_1, t_1, w_2, t_2, f \vDash \Box \forall x((woman(x) \wedge misbehaves(x)) \rightarrow asked\text{-}to\text{-}leave(she_x))$$

iff

for every assignment f′ like f except at most for x, and for every w s.t. $w_1 R w$,

if $S, w, t_1, w_2, t_2, f′ \vDash woman(x)$ and $S, w, t_1, w_2, t_2, f′ \vDash misbehaves(x)$,

then $S, w, t_1, w_2, t_2, f′ \vDash asked\text{-}to\text{-}leave(x)$ and $S, w, t_2, w_2, t_2, f′ \vDash female(x)$

Why some anaphoric uses of pronouns seem to require double indexing while others do not is a question that I shall not attempt to answer.

---

[40] Worse than that, if Tarek is a woman and misbehaves in w_5, that will give us an assignment f′ (f′(x)=Tarek) which makes the antecedent true, but the consequent false, since Tarek is a man at the designated world (alias w_2). So the mere possibility that one and the same individual should be male in some worlds and female in others would make (19) false, which is certaintly not what we want.

1.9. Temporal Indexicals

I now want to go back to the question of how temporal indexicals are analyzed in my account. Insofar as modal indexicals are concerned, there is the indexical 'actually', and that is pretty much it. But in the temporal case, there are many temporal adverbials that may be classified among indexicals. Here are some: 'now', 'then', 'today', 'tomorrow', 'yesterday', 'three days ago', 'last Friday', 'next Thursday', 'in two weeks', 'this evening', 'at 5 pm', etc. The temporal vocabulary is very rich, and most of these words have lexical meanings that are also very rich. Yet, we only saw, back in 1.5.2, how the indexical 'now' is analyzed, namely, with the following truth clause:

$$S, w_1, t_1, w_2, t_2, f \vDash \text{now } \varphi \quad \text{iff} \quad S, w_1, t_2, w_2, t_2, f \vDash \varphi$$

On this definition, 'now' is not vacuous. What it does is to "neutralize" the effects of any temporal operators in whose scope it lies, and to set the time of evaluation back to the designated time. However, though it does not make it a vacuous operator, the proposed clause strips the indexical 'now' of any lexical meaning other than this index-resetting capacity. Now, it is not impossible that the word 'now' is indeed devoid of any interesting lexical meaning. For, what could that meaning be? It might be tempting to suggest that the meaning of 'now' encodes the condition of being the time of some utterance of the word 'now'. However, there are many instances of 'now', for instance in novels, that invalidate this assumption. For example:[41]

(1) He began now to be afraid, whose feelings up to now had been
    bewilderment and perhaps foreboding and fatality.

Also, even though such a condition is not, strictly speaking, a token-reflexive condition (it only requires that there be *some* utterance of the word 'now' at

---

[41] From William Faulkner, *Light In August*. Vintage Books Edition (1987: 296). Such uses of 'now' may be found all over the place. Here, I picked out the book arbitrarily, and opened it arbitrarily, and there I could already find a use of 'now' that I needed.

that time), the resulting analysis may still be vulnerable to certain arguments raised against the token-reflexive theories.[42]

Whatever the case may be with the indexical 'now', there is no doubt that many other temporal indexicals have very rich lexical meanings. But it may not be obvious how double-indexed semantics, in which temporal expressions are treated in a tense logic-style, could account for those. The present section consists of two subsections. In the first one, I will start with some temporal expressions that are (normally) considered non-indexical, such as dates, then move to some "*more* indexical" expressions, such as the names of the days in the week. In analyzing the two, I will also address the issue of the deictic uses of tenses. In the second subsection, I will propose an analysis of the indexical 'tomorrow', which will work, *mutatis mutandis*, for 'yesterday' and 'today'. Then I will look at the expression 'the day after', whose lexical meaning does not appear to be any different from the lexical meaning of 'tomorrow', and try to see what exactly the difference between the two words amounts to.

1.9.1. Dates, Days, and Tenses
Consider the following:

(2) Aaron will get married on Friday, November 9, 2007.

I suggest that we should analyze (2) is such a way that it gets equivalent to the following:

(3) Some day in the future, Aaron will get married, and that day is Friday, November 2, 2007.

The first clause is contributed by the sentence and its future tense, and the second clause is the proper contribution of the temporal adverbial. Ignoring,

---

[42] See Predelli (2006). The argument is, roughly, that on such an analysis, the truth of a sentence such as 'I am hungry now' *logically* entails that the word 'now' exists, or that utterances exist. However, those entailments do not appear to be logical inferences, but merely pragmatic inferences. In 3.4, I will discuss a similar sort of argument, which I will call the argument from unwelcome logical consequence.

for simplicity, the world and the assignment parameters, and taking φ for 'Aaron gets married', the formal analysis of (2) is:

$S, t_1, t_2 \vDash$ sometime-future (Friday-11-2-2007 $\wedge \varphi$)   iff

for some $t > t_1$: $S, t, t_2 \vDash$ Friday-11-2-2007 $\wedge \varphi$ iff

for some $t > t_1$: $S, t, t_2 \vDash \varphi$ and $S, t, t_2 \vDash$ Friday and $S, t, t_2 \vDash$ 11-2-2007

What this means, on the semantic side, is that we will have a bunch of 0-place predicates, such as 'Monday', 'Tuesday' etc. and then the various dates and so on. A structure interprets such predicates, which belong to the non-logical vocabulary, by mappings from times (and worlds) directly into truth values. The interpretation function of 'Friday' will send a time (and world) to True only if the time at stake is a Friday. And so on.

Partee (1973) famously noted that tenses, too, appear to have something like *deictic* uses. Consider:[43]

(4) Aaron will get married.

Tense logic traditionally analyzes (4) as meaning that there is some time in future at which Aaron gets married. But suppose that Aaron's marriage has been scheduled for tomorrow, and it is very important that he should marry, but you are afraid that he might cancel his marriage. If, in such a context, I say "don't worry" and I utter (4), you will take me to mean that he will get married *tomorrow*, and the possibility that he gets married 10 or 20 years from now is, intuitively, irrelevant to the truth of what I am saying in (4).

---

[43] This is not Partee's example, but ours. Partee's example is "I didn't turn off the stove."She points out that the quantifier analysis does not seem to work: if you analyze the sentence as meaning that there is a time in the past at which I did not turn off the stove, then it is trivially true. And if you try something like "It is not the case that there is a time at which I turned off the stove", then it is almost trivially false. It should be noted, however, that quantifiers in general often come with a contextual restriction of their domain. If you restrict the time interval appropriately, you get means of providing a quantifier analysis of those examples, Partee's as well as ours.

If, with Partee, we acknowledge the deictic uses of tenses, how should we analyze (4) in the imagined scenario? It might be tempting to simply propose the following:

S, t, $t_2 \vDash$ Aaron gets married

And then, the extra work that I, *qua* the speaker of (4), and you, *qua* hearer, need to do, is that we pragmatically convene on taking time t, the time of evaluation of "Aaron gets married", to be tomorrow.

The only problem with this analysis is that we would be missing the contribution of the future tense, namely, the fact that time t is posterior (rather than simultaneous or anterior) to time t_2. This problem may be remedied by introducing a new expression in our logical vocabulary, a 0-predicate 'future', defined as follows:

S, t_1, $t_2 \vDash$ future   iff $t_1 > t_2$

Now let us see how this may be put to work.

In the analysis of (2), the temporal existential quantifier in the truth clause for the future tense did not bother us, because the relevant time of evaluation was lexically specified in the sentence. But consider:

(5) Aaron will get married on Friday.

Suppose that I utter this on Wednesday, January 3, 2007. Then you will understand me as saying that Aaron will get married on the coming Friday, which is January 5, 2007. But if we run the same analysis as we did in the case of (2), we get it that (5) is true iff Aaron gets married at some point that lies in the future of January 3, 2007, and that the day of his marriage is a Friday, but it may be *any* Friday, and not necessarily January 5. Those do not seem to be the right truth conditions for (5).

What might work is to say that (5) is true relative to S, t_1 and t_2, iff:

S, t_1, $t_2 \vDash$ Aaron gets married and S, t_1, $t_2 \vDash$ future and S, t_1, $t_2 \vDash$ Friday

And then, we (speaker and audience) need to pragmaticaly convene on taking the intended time of evaluation for (5) to be the 5^{th} of January 2007, rather than some other Friday. In this, we rely, of course, on the context and pragmatics. Semantics only helps us restrain the range of times that could reasonably be intended as the time of evaluation.

1.9.2. 'Tomorrow' and 'The Day After'

When it comes to 'tomorrow', I propose that we introduce into our logical vocabulary a 0-place predicate 'tomorrow', with a truth clause that will go (more or less) as follows:

S, t_1, $t_2 \vDash$ tomorrow   iff t_1 is (on) the day after t_2

For illustration, let's apply it to the following:

(6) Aaron will get married tomorrow.

It will make little difference whether we treat the future tense as being used deictically, or as the existential future modality. That is because 'tomorrow' constrains the interpretation to the day following the day of the designated time, which is determined uniquely (unlike the constrain encoded in 'Friday').

Taking the latter line, we get it that (6) is true relative to S, t_1 and t_2, iff:[44]

for some t such that $t > t_1$:

S, t, $t_2 \vDash$ Aaron gets married and S, t, $t_2 \vDash$ tomorrow

In English and many other languages, the temporal vocabulary contains "minimal pairs" of expressions that do not seem to differ in their lexical meaning, but are not synonymous. One of the two expressions is indexical, with the characteristic feature of indexicals to systematically reset the index of evaluation at the designated index, while the other not only lacks this feature but, it seems, cannot be interpreted as deploying the designated index. So for

---

[44] Here, t_1 will be taken to be the present time, and will normally coincide with t_2. On the deictic construal, the relevant time of evaluation, t_1, will be taken to be the future time at which Aaron is said to be getting married.

instance, one cannot use the following sentence on January 4, 2007, to express that Aaron will get married on January 5, 2007:[45]

(7) Aaron will get married the day after.

'Tomorrow' and 'the day after' form such a minimal pair, and so do 'yesterday' and 'the day before', 'in two days' and 'two days later', 'four weeks ago' and 'four weeks earlier', and so on. The semantic difference in the meaning of the expressions in each pair is best brought out when we consider them embedded:

(8) Maria said that she'd call tomorrow.

(9) Maria said that she'd call the day after.

While (8) reports Maria as having said that she would call on the day after the day of (8), (9) reports her as having said that she would call on the day after she said it. Hence (8) and (9) differ in meaning and in truth conditions.

The existence of such minimal pairs should not be a problem for our account. The semantics of the indexical expression will use double indexing, while the non-indexical expression will deploy only the index of evaluation:

$$S, t_1, t_2 \vDash \texttt{the-day-after}\ (\varphi) \text{ iff } S, t, t_2 \vDash \varphi, \text{ where } t \text{ is the day after } t_1$$

This solves only one part of the problem – namely, it provides the correct truth conditions for (9), but it does not explain why (7) cannot mean that Aaron will get married *tomorrow*.

The question is, then, whether we want the unavailability of those readings, and, more generally, the fact that "the day after" cannot be used for tomorrow and that "five days later" cannot mean "in five days", to be traced to some *semantic* difference, or rather, whether we might look for an explanation that is not necessarily semantic. I opt for the latter. Just as one should not say, for

---

[45] That is, unless the time of evaluation the day after which Aaron is said to be getting married *accidentally* happens to be the day of utterance. For instance, "There will be an earth-quake, and Aaron will get married the day after", as uttered on January 4, 2007, is true if the earth-quake is on that very day and so Aaron gets married on January 5, 2007.

instance, "Maria cares for the children of the woman who has just been mentioned" to mean that Maria cares for *her own* children (cf. 1.6.4), you should not say "I'll call you the day after" to mean "I'll call you tomorrow". It may well be that, in the former as well as in the latter case, the sentences in each pair have the same lexical meanings, hence equivalent semantic contents. The reason why you should not say such things is not, then, that you would be saying something with different truth conditions than those that you want. Rather, it is part of our knowledge of language that such pairs are mutually exclusive, and that whenever you *can* use the indexical, you *ought to* use the indexical. By not doing so, you would be violating certain linguistic conventions – but those need not be reflected in the semantic contents of the expressions at stake.

Let me note that the explanation of the difference between 'tomorrow' and 'the day after', which roughly says that there is no difference between the semantic contents of the two expressions (since their lexical meanings are the same), but that there are different conventions of use associated with the two words, is similar, on the one hand, to the explanation of why the definite description 'the most salient female', while similarly having the same lexical meaning as the pronoun 'she', has a different range of uses (cf. 1.6.4), and, on the other, to a possible explanation of why the negative condition of not being a speaker or a hearer need not be part of the lexical meaning of the 3rd person pronoun. With respect to this last point, recall that simply taking the condition of being female to be the lexical meaning of 'she' was potentially problematic because if the speaker herself is female, she, too, should qualify, it seems, as a possible intended reference. Yet, normally, one should not use a sentence such as "She is hungry" to express that one is hungry.

Some might take this fact to show that it is part of the lexical meaning of 'she' that it should stand neither for the speaker nor for the hearer. My explanation, on the other hand, is that those negative conditions are not part of the lexical meaning, but that the reason why I cannot normally use 'she'

when I am talking of myself is that there are certain linguistic conventions associated with 'she' and 'I' that prevent me from doing so. In particular, there is the convention is that whenever one can use 'I', one ought to use 'I', and not a 3rd person pronoun, or one's name, or a definite description that uniquely obtains of the speaker. Further evidence for this explanation comes from cases in which it does happen that one would use the 3rd person pronoun while talking *de facto* about oneself. Suppose that you would like to meet John Perry, and you go to a party where you hope he would come. Then you start chatting with some friendly guy there, and you ask him:[46]

(10) Is John Perry at the party?

He answers:

(11) Yes, he is.

As you will have guessed, the guy who said (11) is John Perry himself. He simply did not want you to realize that he was John Perry. Now, to bring the point home, observe that (11) is not false. Possibly it is misleading, but it is certainly not false, given that John Perry *is* at the party. However, if we made the condition of not being the speaker part of the lexical meaning of 'he', then (11) would end up being false. (Or, more optimistically, it would lack a truth value – but that is not good enough yet, because (11) is obviously true.)

——

end of Chapter 1

——

---

[46] This sort of example is discussed by Perry in Perry (1997), Perry (2001).

Chapter 2

What Is Said

A major motivation for Kaplan's notion of *content* is that it should stand for the asserted content, or *what is said*. In this chapter, I lay down a series of problems for the existing theories of what is said. However, if we think of semantic content along the lines of the theory laid in Chapter 1, then we may identify semantic content with what is said (and identify both with lexical meaning).

2.1. Introduction

This introduction consists of two subsections. First, I present the standard, Kaplanian view on what is said. In doing so, I will remind you of its main tenets concerning the notions of semantic content and lexical meaning (cf. 1.1), as well as of the distinction between what is said and what is conveyed (cf. 1.4.1). Second, I will tell you what will happen in this chapter.

2.1.1. The Received Wisdom on What Is Said

In everyday life, we say things, and the things that we say may change our lives, affect our relationships and careers, or, to take an extreme case, people may go to jail because of what they have said. There is no doubt that we have a certain intuitive notion of what is said, and attempts have been made to account for it. In philosophy of language, what is said (by someone who utters

a meaningful string of sounds) is, on the standard, Kaplanian view, identified with the *content* of the sentence (relative to its context of utterance). This content, in turn, is seen as the truth-conditional content, that is, as a set of conditions such that the utterance is true if and only if those conditions obtain. Such truth conditions are generally seen as conditions on what the *world* must be like for the utterance to be true (cf. 1.3.2). Suppose that on Friday March 31, 2006, Inma comes into the room, and, pointing at her, I say:

(1) She has graduated in math.

The standard view on what is said concerns all of the following at once:

(i) the semantic value of the pronoun 'she' in (1);

(ii) the semantic content associated with the sentence uttered in (1) (relative to its context of utterance);

(iii) the content that I have asserted by uttering (1), or *what I said*.

Though we have already discussed this view in the previous chapter, let us recall its main tenets:

(i) the semantic value of 'she' in (1) is Inma herself;

(ii) the semantic content of the sentence in (1), relative to the context of utterance of (1), is the proposition that some time before March 31, 2006, Inma graduated in math;

(iii) that same proposition is also the asserted content, or *what is said*.

While the standard view identifies the notion of semantic content with the notion of what is said, it is eager to distinguish the two notions from two related notions: the notion of what is conveyed, on the one hand, and the notion of lexical meaning, on the other.

Consider (1) again, and suppose that I utter it while we are discussing a mathematical problem that we were unable to solve. Then the reason why I told you (1) was not necessarily to *inform* you that Inma has graduated in

math. Let's suppose that you knew that already. Then my intention in (1) must have been to inform you of something else, like the following:

(2) Inma should be able to help us solve the problem.

My utterance of (1) clearly does not *say* the same thing as (2), but in the context at stake, (1) *conveys* what (2) says.

More importantly for our purposes, the standard view draws a sharp distinction between semantic content and lexical meaning. All that the lexical meaning of the sentence in (1) tells you is that prior to some contextually salient time, a female individual graduated in math. Mere lexical knowledge of what the words uttered mean does not enable you to figure out who the woman and what the time at stake are, hence it does not determine the semantic content. However, the move from the lexical meaning to *what is said*, the latter being simply the semantic content, is believed to be fairly direct, requiring only knowledge of some basic parameters of the context of utterance: who is speaking, to whom, where and when, and to what they are referring. The standard view, then, holds that what is said can be determined on the basis of this lexical meaning, the syntax of the sentence, those basic contextual parameters, and nothing else.

On the standard view, both what is said (when the sentence contains indexicals) and what is conveyed depends upon the context. But they do so in different ways. Suppose that on July 24, 2006 at noon I am in Professor Cheng's office, and I say:

(3) It's cold in here.

In the standard view, the semantic content of (3) is the asserted content, and it is the proposition that on July 24, 2006 at noon it is cold in Prof. Cheng's office. This content is obtained more or less directly from the lexical meaning of the sentence uttered: we only need to "plug in" the place of utterance, called for by the indexical 'here', and the time of utterance, called for by the present

tense. By contrast, what is conveyed depends heavily on the context. Typically, an utterance of (3) will convey something like the following:

(4) I would like you to close the windows.

However, suppose that July 24 is a very warm day, and that Prof. Cheng's office is so air-conditioned that it is freezing in there. Then my utterance of (3) may convey exactly the opposite of what it conveyed in the previous scenario:

(5) I would like you to *open* the windows.

Hence figuring out what is conveyed requires reasoning about speaker's beliefs and intentions, inference to the best explanation, and a fair amount of encyclopedic knowledge.

The ways in which *what is said* and *what is conveyed* depend on the context has led some some philosophers, like Kent Bach, to distinguish between the notions of *narrow* and *broad* context. Caricaturing somewhat, the standard view may thus be captured with these two equations:

– linguistic meaning + narrow context = what is said

– what is said + broad context = what is conveyed

2.1.2. A Short Preview of What Follows

I have two goals in this chapter. First, I will present a series of cases that put the mainstream view into jeopardy. What those cases show is that neither our practices of reporting what is said nor our intuitions on what is said fit into the Kaplanian model. Second, I will suggest that those cases fit better into a simpler, single-level model, in which the notions of semantic content and what is said are still interchangeable, but are both identified with the notion of lexical meaning. In other words, my view fits into this much simpler equation:

– lexical meaning (=what is said) + context = what is conveyed

In rough lines, in the model that I am going to propose, what is said in (1) ("She has graduated in math") is merely, on a first approximation, that prior to some contextually salient time, some female graduated in math. The semantic value of the pronoun 'she' in (1) is not Inma herself, but the general condition of being female, lexically encoded in the meaning of the pronoun. Granted, there is the intuition that when I utter (1) pointing at Inma, I say something *about* Inma. This intuition is preserved through the independently motivated notion of *that of which/whom the speaker is talking* (cf. 1.2). The idea is that when I utter (1), I am talking about Inma and am, *qua* speaker, referring to Inma. However, Inma is not part of what I said – rather, what I said is the lexical meaning of the sentence (which is also the semantic content, as argued in Chapter 1), but I assert that meaning *of, or about,* Inma. After laying down my view, I will end the Chapter by reconsidering, one by one, the cases that I will have previously shown to be problematic for the standard view.

2.2. Two Motivations for the Mainstream View

At a first glance, one might plausibly suppose that the meaning lexically associated with a sentence is the most obvious candidate to play the role of *what is said* by an utterance of that sentence. I am going to argue that this is basically correct. However, this idea is widely rejected nowadays. It will help, then, to start with those cases that have motivated its rejection.

2.2.1. Different Meanings, Same Things Said

The first motivation for a distinguished level of what is said comes from utterances whose speakers intuitively say the same thing, even though the sentences that they use do not have the same lexical meaning. Suppose that I say:

(1) I have graduated in philosophy.

By uttering (1), I may inform you that I have graduated in philosophy. Now, suppose that you want to inform someone else of this. You cannot use the same sentence that I used, because then you would inform your interlocutor that you, not I, have graduated in philosophy. I can sure refer to myself using the first person pronoun, but you need to find another way of referring to me. For instance, you might say:

(2) Isidora has graduated in philosophy.

Or, if I am there and you are pointing at me, you might just say:

(3) She has graduated in philosophy.

The sentences uttered in (1), (2) and (3) have different meanings, given that different lexical conditions are associated with 'I', 'she' and proper names. The 1st person pronoun is used for the speaker, while the 3rd person pronoun, for some salient female. The name, if there is any condition lexically associated with it at all, will just tell you that it stands for its bearer. Still, for Kaplan and his followers, what you say in (2) or (3) and what I say in (1) is one and the same thing – something like the proposition true in those and only those worlds in which I, Isidora, have graduated in philosophy before the time of my utterance. This intuition goes back at least to Frege, who wrote:

> It is not necessary that the person who feels cold should himself give utterance to the thought that he feels cold. Another person can do this by using a name to designate the one who feels cold. (1899: 236)

2.2.2. Same Meanings, Different Things Said

The second motivation for the dominant view is the idea that you can use one and the same non-ambiguous sentence to express different things, provided that you use it in different contexts. As Kaplan puts it:[47]

---

[47] It should be noted that this insight, too, goes back to Frege: "The sentence 'I am cold' expresses a different thought in the mouth of one person from what it expresses in the mouth of another." (*ibid.*)

> What is said in using a given indexical in different contexts may be different. Thus if I say, today, "I was insulted yesterday," and you utter the same words tomorrow, what is said is different […] There are possible circumstances in which what I said would be true but what you said would be false. Thus we say different things. (1977: 500)

Since Kaplan, it has been widely held that once we have indexicals in the language, lexical meaning differs from semantic content, and thereby from what is said, in two respects: (1) there is something in what is said that is not in the lexical meaning, namely, the reference of indexicals; (2) there is something in the lexical meaning that does not reach into what is said, namely, the lexically encoded conditions that merely help fixing the reference, such as being the speaker for the indexical 'I', or being female for the pronoun 'she'.

2.3. The Theoretical Landscape Today

It is interesting that almost everyone today accepts this twofold difference between lexical meaning and what is said. Almost everyone holds that the referent of an indexical expression is a constituent of what is said without being part of the lexical meaning, while the general conditions lexically associated with the indexical are part of the lexical meaning without reaching into the semantic content and into what is said. This view on the indexicals' contribution to what is said appears to be so entrenched in philosophy of language that it is sometimes the only point of agreement among authors whose views are otherwise radically different. However, what makes the current debate more complicated is the lack of agreement on the issue of whether what is said, or the asserted content, is the semantic content, and if not, on the question of how the two notions are related to each other.

Rather than try to depict the theoretical landscape in a thorough and systematic way – a task that would by itself require a book – I would simply like to take a few glances at this landscape, so that we may be aware of the

variety of approaches to the notions of semantic content and of what is said defended today.

2.3.1. The Main Camp (Literalism)

The view that has emerged from Kaplan's *Demonstratives* is still, I believe, the most widespread view. Although there may be many differences and points of disagreement among the philosophers who endorse this view, they all draw a sharp distinction between linguistic meaning and what is said. Furthermore, they take what is said to be propositions, traditionally conceived; which is to say, *eternal* propositions, whose truth values depend on what the world is like, but do not vary with other things, such as times or places or individuals.[48]

Literalists hold that in the case of "pure" indexicals, such as the 1st and the 2nd person pronouns and adverbs 'here', 'now', 'today' and the like, the linguistic meaning itself is determinate enough to fix the reference, and hence determine what is said by the utterance. For example, the linguistic meaning of 'I' tells you that its referent is the speaker, and since there is always a unique speaker in any given context, the move from the lexical meaning to what is said is pretty much automatic. But they also accept that in certain cases, the linguistic meaning is too poor to determine what is said. Demonstrative pronouns are the case at point: the linguistic meaning of 'he', which boils down to something like 'salient male', only constrains the referent of 'he', but in order to determine the actual referent, who will then be part of what is said, one must take into account certain pragmatic facts about the utterance, such as which individual was jointly attended to by the speaker and her audience in the context of utterance. However, the appeal to pragmatics is triggered here by something from the sentence itself, namely the 3rd person pronoun. The

---

[48] The main camp is divided between those who conceive of propositions as structured entities and those who model propositions by sets of possible worlds. It should be noted, too, that in his formal system, Kaplan models contents as functions that map pairs consisting of a possible world and a time to truth values, which is to say that such contents can and normally do take different truth values at different times. However the informal understanding of what is said that has subsequently emerged from Kaplan's work remains tied to the traditional notion of proposition, that is, something that, if true, is true once and for all, and if false, is similarly false once and for all.

trademark of literalism is precisely that all the pragmatics involved in the determination of what is said must be constrained by semantics.

Most literalists acknowledge that there need not always be an overt, phonetically articulated element in the sentence to trigger the pragmatic resolution of an element that will go into what is said. Consider:

(1) Every bottle is empty.

If Sonia utters (1) at a party, she likely means to be saying that every bottle *at that party* is empty. In many views, the intended domain of quantification is taken to be part of what is said. For instance, what is said by (1) would be the proposition true in those and only those worlds in which every bottle at the party is empty (at that world, at the time of utterance). In the context described, we take (1) to be false if some bottle at the party is not empty, but we do not take (1) to be false just because there is some bottle somewhere in the world that is not empty. If people's intuitions on truth values are to be taken seriously, and if what is said is conceived of as being a proposition, which varies in truth value along the possible world dimension, but not along other dimensions such as times, locations or domains of quantification, then it is easy to understand why the domain of quantification and other parameters that similarly affect the truth value have been seen as elements of what is said.

Indexicalism, a variant of literalism, holds that for any element that reaches into what is said, there must be some element in the *syntax* of the sentence whose semantic interpretation will trigger a pragmatic process, such as those involved in interpreting demonstratives, that will help determine what is said. As Stanley writes:

> All effects of extra-linguistic context are traceable to elements in the actual
> syntactic structure of the sentence uttered. (2000: 391)

Be that as it may, what literalists insist on is that the pragmatic processes that bridge the gap between the lexical meaning and the semantic content, thereby

bridging the gap between the lexical meaning and what is said, should always be monitored by syntax or semantics.

2.3.2. *Contextualism*

It is on this very last point that contextualism departs from literalism. Contextualists hold that the lexical meaning, to the extent that the notion makes sense at all, only loosely constrains what is said. Various pragmatic processes enter the determination of what is said without being triggered by any syntactic or semantic property of any element in the sentence. Consider:

(2) Inma looked into her purse, found the key, and opened the door.

A contextualist wants to say that part of what is said in (2), in the situation imagined, is that Inma found the key *in her purse* and that she opened the door *with that key*. But there does not seem to be any element in the sentence uttered that invites the interpreter to identify Inma's purse as the place where she found the key, and there is nothing that stands for how she opened the door (viz. with the key). From all that is syntactically and lexically encoded in the sentence, (2) should be true even if Inma found the key lying on the floor and opened the door without using the key. However, for the contextualist, such "minimal" truth conditions are not empirically adequate. They do not capture the intuitive truth conditions that a competent speaker would associate with an utterance of (2) in the imagined scenario. The lesson that they draw is that syntax and semantics provide only some rough guidelines to computing the truth conditions and to determining what is said. Often, elements that come directly from the conversational context may "intrude" into what is said. It has thus become customary to talk of *pragmatic intrusion.* A particular instance of this, which we find in the case of (2), is what Recanati calls *free enrichment* (see Recanati (1993: 258-260), (2004: 23-28)).

Unfortunately, there is quite a lot of terminological confusion when people talk about contextualism. To begin with, "contextualism" means something different in current epistemology than it does in the semantics/pragmatics

debate. A contextualist in epistemology holds that the verb 'know' depends on the context, but this context-dependency is entirely compatible with *lexically encoded* context-dependency, hence such a "contextualist" need not be a contextualist in the sense relevant to the present discussion.

More annoying is the fact that even in the semantics/pragmatics debate, different authors mean different things by "contextualism". For instance, Jason Stanley or Kenneth Taylor, whom I would take to belong in the literalist camp, are considered by Cappelen and Lepore (2005) as "moderate" contextualists. Recanati (2004) proposes a somewhat different taxonomy. He reserves the contextualist label for those views that allow optional pragmatic processes, not triggered by anything in the sentence, to participate in the determination of what is said. Contextualism, as Recanati understands it, holds that the "minimal" proposition determined by the syntactic and semantic properties of the sentence uttered, supposing that there is such a thing, does not play any role in determining what is said. It is also in this sense that Predelli (2006) uses the term 'contextualism'. Typical contextualists are people like Charles Travis and John Searle, and their claim is that pragmatics can intrude into semantic content is a way that is lexically, syntactically and semantically *un*constrained. In this work, I am adopting the Predelli-Recanati taxonomy.

2.3.3. *John Perry's Reflexive-Referential Theory*
In his recent work, John Perry has been arguing that it is misleading to talk of *the* proposition expressed by an utterance or of *the* truth-conditional content. For, there is a wide array of propositions that may be associated with any given utterance, all of which provide, in one way or another, a necessary and sufficient condition for the utterance to be true. Suppose that on July 15, 2005, David Kaplan says "I was insulted yesterday." The semantic or truth-conditional content that the standard view associates with this utterance is the proposition that Kaplan was insulted on July 14, 2005, while in Perry's view, that is simply one among the contents associated with the utterance. There is also what he calls the *reflexive* content (or, one among the reflexive contents),

which captures the truth conditions that any competent speaker is able to associate with the utterance, independently of any further knowledge about the context of utterance. Thus, in virtue of what the words uttered by Kaplan mean in English, we know that his utterance is true if and only if the speaker of that utterance was insulted on the day before the utterance. For Perry, lexical meaning is thus captured at the level of reflexive content.

In Perry's theory, the move from this reflexive content to the Kaplanian truth-conditional content normally consists in several steps, each of which results in a new content that specifies a certain truth condition for the utterance. For example, if you start from the reflexive content of Kaplan's utterance above, and you add the fact that it is David Kaplan who made it, you obtain that the utterance is true if and only if *Kaplan* was insulted on the day before the utterance. Similarly, if you start from the same reflexive content and add the fact that the utterance was made on July 15, 2005, obtain get that the utterance is true if and only if the speaker, whoever he or she may be, was insulted *on July 14, 2005*. Finally, if you "increment" the reflexive content with the two facts together, you obtain that the utterance is true iff David Kaplan was insulted on July 14, 2005, which is the Kaplanian content associated with the sentence that Kaplan uttered.

Having stressed that "the binary distinction (...) is too simple," Perry thus writes:

> An utterance has as wide a variety of contents as we may find useful to isolate, for particular purposes of description and explanation. We can say that in at least the vast majority of cases, the common sense concept of "what is said" corresponds to content$_C$ [= the usual propositional content]. This is a good reason for an account of content to recognize this concept, but not a good reason to expect it to be the only or even the most theoretically fruitful kind of content" (1997: 17)

In sum, although Perry's view lines up with the traditional view insofar as the notion of what is said goes, it is more flexible and, one may hope, better equipped to deal with certain problematic cases.[49]

2.3.4. Robert Stalnaker's Theory

Stalnaker's view is very often assimilated to Kaplan's view, but there are some important differences, which I will try to point out in this subsection. Stalnaker is famous for having proposed a theory that uses two-dimensional matrices, which he also calls *propositional concepts*, and those are simply propositional functions: functions that take as arguments things along the one dimension and return propositions, which are, in turn, functions from the things along the other dimension to truth values. For Stalnaker, both dimensions consists of possible worlds. His overall picture thus resembles Kaplan's, if you think of the worlds along the one dimension as Kaplanian contexts, and the worlds along the other dimension as circumstances of evaluation. In turn, propositional concepts come pretty close to Kaplanian characters.

Before I draw your attention to some aspects of Stalnaker's theory that show it to be actually quite different from Kaplan's, let us just make sure that we are clear on what Stalnaker's matrices looks like and how they work. Consider the following sentence:

(3) I am a famous singer.

Now consider four possible worlds, α, β, γ, and δ, that are as follows. In α and β, the sentence in (3) is uttered by Toto Cutugno, while in γ and δ, the same sentence is uttered by Stefano Predelli. In α and γ, Toto is a indeed a famous singer, while in β and δ, he is something else (say, a pizzaïolo). In all the four

---

[49] In Stojanovic (2003), I have shown how Perry's theory can handle certain cases that Kaplan's theory cannot. But I also pointed out that Perry's theory, as it stands, is too flexible, and that there are utterances for which it would predict that they express the same content while, intuitively, they do not.

worlds, Predelli is a philosopher. The propositional matrix that we get for (3), with respect to this set of worlds, is as follows:

	α	β	γ	δ
α	T	F	T	F
β	T	F	T	F
γ	F	F	F	F
δ	F	F	F	F

Table 1.

Note that the propositional matrix in *Table 1* does not correspond yet to a full-fledged Kaplanian character. That is because the matrix is construed over a limited set of words. So, for example, the proposition that (3) expresses with respect to worlds γ and δ is, according to this matrix, a *necessarily false* proposition: throughout the whole row, it only gets the value False. But of course, one would not want to say that it is necessary that Predelli should not be a famous singer. He might well have been a singer rather than philosopher, and might well have been famous as a singer. To get to a full-fledged character, you would need to construe the propositional matrix over the entire set of *all* the possible worlds.

Stalnakerian propositional concepts are typically construed upon some *restricted* set of possible worlds, namely, upon what he calls the *context set*. The context set consists of the possible worlds all of which satisfy the propositions whose truth is taken for granted by the speaker and the hearer, and are such that both the speaker and the hearer know that the other is taking their truth for granted, and so on. In other words, the context set consists of the worlds which correspond to the speaker's and the hearer's background knowledge. For example, if I we have been talking to each other in English, those possible worlds in which either of us does not speak any English – worlds that are

metaphysically certainly possible (e.g. I might have never learned any English) – will be excluded from the context set. Or, if we both know that Paris is the capital of France, no possible world in which it isn't will figure in the context set. So, the context set is relatively narrowed, but it still contains an outrageous number of possible worlds: all the issues that have not been settled, perhaps because they belong to the future (such as the issue of whether it will rain tomorrow), or because the speaker and the hearer have different opinions on the issue (such as the issue whether Inma is smart) or because neither of them has ever thought about the issue (such as the issue of whether David Kaplan had eggs for breakfast on April 23, 1973) will keep generating possible worlds that will all belong the context set.

This is a good place to introduce an idea, due to Stalnaker and Lewis, that has given rise to a branch of dynamic semantics known as *update* semantics. The idea is that information consists in the elimination of epistemic alternatives that were up to then "live options". Let me explain. Suppose that you meet Toto Cutugno, he tells you his name, but you have actually never heard of him before. You ask him what he does and he utters the sentence in (3), repeated below:

(3) I am a famous singer.

Here, your "epistemic set" (that is, the set of possible worlds that, for all you know or believe, might turn out to be the actual world) includes only the worlds in which (3) is uttered by Toto, but in some of those worlds he is a bank employee, in some he is a tailor, in some he is a singer, and so on. Simplifying considerably, your epistemic set – which will actually coincide with the context set – reduces to the worlds α and β previously used. The propositional matrix that corresponds to (3) over that pair of worlds is simply:

	α	β
α	T	F

	α	β
β	T	F

Table 2.

The upshot of Toto's utterance of (3) is, then, to enable you to narrow down the set of your epistemic alternatives, which you do by eliminating all those worlds in which the proposition expressed is false (assuming, of course, that you assent to the truth of his utterance). In our example, you would simply drop the world β – the world in which Toto is not a singer.

So far we could observe (at least) three things with respect to which Stalnaker's theory is quite different from Kaplan's. First, while Kaplanian contexts are quadruples (agent, place, time, world) and circumstances of evaluation are world-time pairs, Stalnaker uses possible worlds in both roles. I will not discuss this difference. Secondly, Stalnaker has the notion of *context set*, and in fact, he thinks that the best way of formally capturing the informal notion of *context* is precisely as the set of possible worlds that constitute the conversational background – the worlds that satisfy the propositions whose truth is taken for granted by the speaker and her audience (Stalnaker (1998)). Thirdly, we saw how the notion of *update*, or of the elimination of possible worlds in which the proposition expressed is false, plays an explanatory role in Stalnaker's account of information.

There is yet a fourth striking difference. Inspired by Grice, Stalnaker (1978) proposes three "principles" that govern conversation and can be formally accounted for within his framework:

> 1. A proposition asserted is always true in some but not all of the
> possible worlds in the context set.
>
> 2. Any assertive utterance should express a proposition, relative to each
> possible world in the context set, and that proposition should have a
> truth value in each possible world in the context set.

> 3. The same proposition is expressed relative to each possible world
> in the context set.

Principle 1 roughly says that the asserted proposition should be neither trivially true nor trivially false. It combines the ideas from Grice's maxims of quantity and quality. That the asserted proposition should not be true in all the possible worlds means that it should be informative, i.e. that you need to be able to eliminate some worlds (cf. Quantity). That it should be true in some world means that you should not assert what you presuppose to be false (cf. Quality). Principle 2 tells you, for example, not to use empty terms, such as a name without a bearer, or a demonstrative without a demonstratum (because no proposition will get expressed), and not to express propositions that have truth value gaps with respect to the context set. For instance, if it is not presupposed that Tarek had been drinking, you should not say that he has stopped drinking, because the proposition expressed will lack truth value in the worlds in which he had not been drinking. Finally, Principle 3 tells you to avoid ambiguity (but we will shortly discuss this principle in greater detail) .

When a principle is violated, there are basically two repair strategies. One is to conclude that the context set is different from what one thought it was. The other is to conclude that what is said, or the asserted proposition, is different from what one thought it was. And one way in which the second strategy can be most easily made to work, according to Stalnaker, is by projecting the "diagonal" proposition over the matrix, so that it is actually the diagonal proposition that gets expressed relative to every world along the first (vertical) dimension.

Let us see how this works. Suppose that you go to a party, and you know that there will two Italians there, Toto Cutugno and Stefano Predelli. You have never met them before, but you know that Predelli is a philosopher, though you don't know that Cutugno is a singer. Then you meet this guy who looks very Italian and speaks with a bit of an Italian accent and who tells you (3). Note that your epistemic set is now exactly the set {α, β, γ, δ} that we had in

Table 1. Now go back and look at the matrix in *Table 1*. You will see that it does not conform to the conversational principles laid down earlier. First, not the same proposition is expressed in the four worlds. Second, the proposition expressed in γ and δ violates Principle 1, as it is always false.

There are several repair strategies. Given that the presence of γ and δ in the context set violates Principle 1, you could conclude that the context set was α and β, and that Toto was assuming that you knew who he was. But that need not be the case: if you never met Toto before, why would he presuppose that you know who he is. So, following the typical repair strategy when Principle 3 is violated, you could consider that what he said – the proposition that he actually asserted – is the diagonal proposition. So then the propositional matrix over those four worlds becomes as follows:

	α	β	γ	δ
α	T	F	F	F
β	T	F	F	F
γ	T	F	F	F
δ	T	F	F	F

Table 3.
(the matrix of the diagonal
proposition from Table 1)

Having made the original matrix conform to the three principles, you may proceed to the usual sort of update: eliminate all the worlds in which the proposition is false, which leaves you with the world α.

Let me close with two quick remarks. First, the previous example is also an example of how the usual factual knowledge, as your knowledge that Stefano Predelli is a philosopher, enables you to eliminate uncertainties as to who was the speaker of (3). For, by eliminating the worlds γ and δ in the matrix of *Table*

3, you have eliminated the worlds in which Predelli is the speaker of (3), and have thereby learned that Cutugno was the speaker of (3). Note how different this is from Kaplan's picture, in which (3) simply expresses the singular proposition that Toto Cutugno is a philosopher, but to see that this is the proposition expressed, and that the speaker is Toto Cutugno, you would need to engage in a meta-semantic reasoning, such as asking yourself what would have been the proposition expressed if this and such and so.

The second remark, more relevant to the topic of this chapter, concerns the indexicals' contribution to what is said. Stalnaker's proposal departs from the mainstream referentialism in one important respect. If the speaker does not presuppose that the hearer knows who the speaker is, then the proposition expressed will not be a singular proposition involving the speaker. Consider *Table 3*. If Toto does not presuppose that you know who he is, and he simply tells you (3), then the proposition asserted is the diagonal proposition. So, in the example under consideration (with the context set {α, β, γ, δ}), what is said by Toto using the 1st person pronoun is equivalent to the following:

(4) Either Toto Cutugno or Stefano Predelli is a famous singer.

Though this may seem odd and counter-intuitive at a first glance, it is actually a good thing about Stalnaker's proposal. It makes the notion of what is said more flexible, and might be able to handle certain cases that will be shown to be problematic for the mainstream view.

2.3.5. Kent Bach's Theory of Implic*i*tures

The notion of implic*i*ture – to be distinguished from the Gricean notion of implic*a*ture –, was introduced by Kent Bach. There are two motivations for this notion. First, suppose that Sonia says:

(5) Olaf is ready.

What is, one might ask, the truth-conditional content of Sonia's utterance of (5)? It seems difficult to hold that her utterance is true if and only if Olaf is

ready, because it is not clear what the world should be like for Olaf to be ready *simpliciter*. Rather, to be ready is to be ready *for* something, and that for which Olaf is said to be ready is something that the hearer must figure out from the context of utterance. Now, in Bach's view, what is said by Sonia is indeed nothing more than that Olaf is ready, but he also points out that this is not a full-fledged proposition. It is what Bach calls a propositional radical, which can be fleshed out into a proposition by specifying that for which Olaf is said to be ready. For instance, what Sonia might want to communicate is that Olaf is ready to leave home to go to work. This process Bach calls *completion*, and what you obtain is the proposition that Olaf is ready to leave home to go to work, which is the *impliciture* associated with Sonia's utterance of (5). It is the impliciture that gets communicated, even though what is said is only the propositional radical.

The second motivation for implicitures comes from those cases in which there is a mismatch between what the speaker literally says and what it is reasonable to suppose that she means to be saying. Reconsider (1), as uttered by Sonia at the party:

(1) Every bottle is empty.

Here, what is said, according to Bach, is already fully propositional. There are determinate truth conditions for her utterance, namely, it is true if and only if every bottle (simpliciter) is empty. However, as noted above, it is more reasonable to suppose that Sonia means to be saying only of the bottles at the party that they are empty. This process Bach calls *expansion*, and it, too, gives rise to implicitures, such as the proposition that every bottle *at the party* is empty:

> Impliciture can be a matter of either filling in or fleshing out what is said.
> Completion is the filling in of a propositional radical, and expansion is the
> fleshing out of the minimal proposition expressible by an utterance. (1994:
> 144)

What distinguishes implicitures from the usual Gricean implic*a*tures is that they are "closely related to what is said" (126). Impliciture is what gets communicated but is already *implicit* in what is said. Implicature, on the other hand, is something that the hearer must infer, using general knowledge, broad context, and inference to the best explanation, from what the speaker said or communicated, together with the fact that she said it.

It should be stressed that Bach insists that his notion of what is said is a technical notion and cares little about speakers' intuitions on what is said. He also thinks that discourse reports are irrelevant to the notion of what is said that he is interested in. But this is precisely why his theory, as a theory of what is said, lacks interest. For, if you posit some notion as a technical notion, you can build upon it any theory you like, so long as it remains coherent. The theory will only be of interest if it sheds light on some interesting empirical phenomenon. For instance, a theory that uses a technical notion of what is said is worth pursuing if it can illuminate the intuitive notion of what is said, or account for the use of the locution 'what is said', or if it can be worked into a theory of indirect discourse reports, or, ideally, if it can do all of that. In the later parts of this chapter, when I set out to propose a theory of what is said in terms of lexical meaning, those will be the driving motivations for my theory. So, I take it to be obvious that a theory of what is said should strive to be able to account for people's intuitions on what is said, intuitions which can be hardly disentangled from our practices of discourse reporting. Hence, even if we grant Bach the technical status of his notion of what is said, what remains to be seen is whether his notion plays any interesting role in an account of the intuitive notion of what is said. Now, to be sure, Bach defines impliciture in terms of completion and expansion of what is said. But what he fails to show is that those processes could not operate directly on the notion of lexical

meaning. It thus remains unclear whether the intermediary, technical notion of what is said is really doing any interesting work in Bach's general proposal.[50]

2.3.6. Semantic Minimalism

Semantic minimalism, a view defended in Borg (2004) and Cappelen and Lepore (2005), is a variant of literalism that wants to minimize, if not outright abolish, the role of pragmatics in the determination of semantic content and of truth. Reconsider once more Sonia's utterance at the party:

(1) Every bottle is empty.

Minimalists would deny that the domain of quantification relevant to the truth value of (1) is in any way dependent upon the context. They would insist that (1) is true if and only if *every bottle is empty*, so that some non-empty bottle somewhere in the world suffices to falsify (1), regardless of how things stand at the party.

Oddly enough, when it comes to indexicals, minimalists happily accept that the semantic content is richer than the lexical meaning, in that it includes the contextually specified referents of such expressions, though not the lexical meaning itself. On this score, there is hardly any divergence between contextualism and minimalism, views believed to be radically opposed.

It should be stressed, however, that minimalism about semantic content is not *ipso facto* minimalism about what is said. Indeed, Borg, Cappelen, and Lepore all reject the assumption that the semantic content is also the asserted content, or what is said. Cappelen and Lepore thus believe in "speech-act pluralism" and hold that while the semantic content associated with an utterance is determined by the lexical meaning of the sentence and only the most basic contextual factors, such as who is speaking, where and when, there are indefinitely many things "said" by any given utterance, because there can be indefinitely many intuitively correct reports of what is said, which need not

---

[50] Robyn Carston (forthcoming) convincingly argues that Bach's level of what is said, as distinct from lexical meaning, is not well-motivated.

have any content in common. Borg also thinks that speech reports are relevant to the notion of what is said, and accepts that pragmatic factors may affect what is said, but she thinks that semantic content and what is said are notions that have nothing to do with one another.

2.3.7. David Lewis's Theory

All of the views that we have seen so far, however they might diverge in various respects, appear to share some core assumptions regarding the semantic content of indexical sentences. More or less all of them endorse Kaplan's view and agree on the following:

(i) there is something in the semantic content that is not in the lexical meaning, namely, the reference of indexicals;

(ii) there is something in the lexical meaning that does not reach into the semantic content, namely, the lexically encoded conditions that merely help fixing the reference, such as the condition of speakerhood for the indexical 'I'.

A notable exception is David Lewis. As already noted in 1.5.5, Lewis was careful to point out that mere double-indexed semantics (even with distinct types of indices) does not commit one to holding that semantic content should be the Kaplanian content, rather than the character. Lewis thus writes:

> Let us agree that sentences depend for their truth on both context and index. What then, should we take as their semantic value? (1980: 92)

He then proposes two options, the first of which is, roughly, to take Kaplanian contents as semantic values, and here is the second:

> The semantic values of sentences are constant but complicated. A value for a sentence is a function, perhaps partial, from combinations of a context and an index to truth-values. (*Ibid.*)

Such "constant but complicated" semantic values come very close to lexical meanings. Lewis then goes on to argue that there is no compelling reason not

to think of semantic value (or what I have been calling semantic *content*) as being constant functions from context-index pairs to truth values.

Now, Lewis was one of the first to express doubts on the very enterprise of providing a *semantic* account of what is said, in which he was followed by a fair number of philosophers and semanticists.[51] He famously wrote:

> Unless we give it some special technical meaning, the locution 'what is said' is very far from univocal. It can mean the propositional content, in Stalnaker's sense (horizontal or diagonal). It can mean the exact words. I suspect that it can mean almost anything in between" (1980: 97)

Let us not be sloppy in interpreting this passage. What Lewis remarks there is that the locution 'what is said' is ambiguous. From this remark to the observation that the *notion* of what is said is vague or obscure or not well-defined, the transition relies on the assumption that our only tool in studying the notion of what is said is through the use of the locution 'what is said' (and locutions of the same ilk). I believe that Lewis makes this assumption, and I believe that this *is* a plausible assumption indeed. I thus think that it is fair to take Lewis to be holding the view that because the notion of what is said is too versatile, we should not hope to get a grip on it within our *semantic* theory.

I believe that Lewis and all those who have been eager to give up the idea that semantic content is also the *asserted* content have despaired too quickly. It is true that the locution 'what is said' is "far from univocal". But it does not follow that the ambiguity of this locution escapes every semantic or even syntactic theorizing. I will try to show that the locution 'what is said' follows patterns that allow for semantic modeling, and that it is therefore possible to get some more or less robust grip on the intuitive notion of what is said. What is bad news for Kaplan is that those patterns do not fit into the schema "what is said =(Kaplanian) content". We are also going to see (e.g. in 2.5) that the use of the locution 'what is said' is not as unconstrained as Lewis seems to think, and that he is thus wrong to suspect that "it can mean almost anything in

---

[51] For example, I. Heim, S. Predelli, N. Salmon, S. Soames, the minimalists, and so on.

between". Finally, another reason that those who sought to banish the notion of what is said into the realm of pragmatics might have had for doing so is that there is a range of cases where the locution 'what is said' is indeed used in ways that seem to grow out of control. Those are the cases that fall under what Paul Ziff (1972) called "the implication sense of what is said." However, I am going to show (in 2.6) there is a sharp delineation between such cases that involve a *loose* use of the locution 'what is said', and the cases that involve a *literal* use of this locution. Again, the bad news for the standard view is that the literal use itself does not get captured by the Kaplanian schema.

2.4. *What Is Said* and Time

In surveying the theoretical landscape regarding the notion of what is said, we pointed out, albeit roughly, the differences among the various view. Thus, for instance, both Perry and Stalnaker remain fairly faithful to the Kaplanian picture, while giving their theories some new twists, so that they can deal with a greater range of problems. On the other hand, a view like contextualism presents a radical departure from the mainstream view, since it bids us to give up the very enterprise of truth-conditional semantics. However great the differences among the views that we have surveyed, we have also seen that most of them agree on the issue of what it is that indexicals contribute to what is said, and that it is their reference. True enough, a number of philosophers and linguists (e.g. Lewis, Heim, Predelli, Salmon, Soames, the minimalists) think that what is said, or the asserted content, is not necessarily the semantic content. In section 2.6, I will discuss some cases that might have led those philosophers to despair that there could be any theory of what is said that identifies this notion with the semantic content, and I will try to show that the reasons for despair are not good enough. In the meantime, in this and the next section, I am going to present cases that are problems for the mainstream view, and for pretty much all the views that agree on the assumption that indexical

contribute their reference, and nothing but their reference, to what is said. I will start with some questions regarding the contribution of tenses to what is said.

2.4.1. How Time Affects What Is Said

It is widely held that in simple sentences, the present tense works like a referential expression (cf. Partee (1973)). It picks out a time, presumably the time of the utterance, and brings it into the semantic content and into what is said by the utterance. The problem is that if what is said is thus tied to a specific time, there will be many cases in which people have the intuition that the same thing has been said, and will easily report what has been said as being the same, even though the propositions expressed by the reported utterances do not coincide on the time picked out by the present tense. To see the point, it is enough to go back to one of our previous examples:

(1) She has graduated in math.
   (me, talking of Inma, on Friday, March 31, 2006)

(2) Inma has graduated in math.
   (Tarek, talking of Inma, on Tuesday, April 3, 2006).

(3) That's what Isidora said, too.
   (a possible reply to Tarek's utterance of (2))

The reply in (3) is intuitively correct. It is natural to take Tarek to have said, in (2), the same thing as I did in (1), even though we spoke on different days, and therefore our utterances have different Kaplanian contents, namely, that Inma has graduated prior to March 31, 2006, the time of my utterance, vs. that she has graduated prior to April 3, 2006, the time of Tarek's utterance. *En passant*, note that not only the contents, but the lexical meanings of the sentences uttered are different, too.

Some might think that the reason why we can so easily that I said the same thing as Tarek is this. It is a fact that if some event e has happened prior to

time t_1, and if t_1 is before t_2, then event e has happened prior to time t_2. It follows that the truth of my utterance entails the truth of Tarek's utterance. Now, whether or not this is the correct explanation of the intuitive same-saying in (1)-(2), where the tense of the sentence is present perfect, it is easy to devise sentences in present progressive that exhibit the same pattern, but for which the explanation does not work. Consider:

(4) Inma is writing a paper on Montague.
   (said by Tarek, on March 31, 06)

(5) I am writing a paper on Montague. (said by Inma, on April 3, 06)

(6) That's what Tarek told me.
   (a possible reply to Inma's utterance of (5))

Again, the reply in (6), as a report of what Tarek said in (4), is correct, though no explanation in terms of the truth of (4) entailing the truth of (5) seems available.

2.4.2. Writing a Paper vs. Having Dinner

Of course, there may be other explanations available to the defenders of the dominant view. For instance, one could point out that writing a paper is an action that normally takes several days, even months, so that the time interval picked out by the present progressive in (4) will sufficiently overlap with the time interval picked out by the tense in (5), so that even if the propositions expressed by (4) and (5) are not exactly the same, they will be similar enough to be reported as being the same.

Variations on the example seem to provide evidence that supports this explanation. For, suppose that Inma utters (5) at a time very distant from the time at which Tarek utters (4): suppose she says it 5 years from now (i.e. in 2011). Then it does not seem correct to reply (6) on the grounds of Tarek's utterance of (4). Or, at least, some qualification would be needed, such as "That's what Tarek told me *five years ago*."

To get to the same point, consider what happens in reports of what is said when the reported action concerns a short period of time:

(7) I am having dinner. (said by Inma, on March 31, at 7 pm)

(8) Inma is having dinner. (said by Tarek, on March 31, at 11 pm)

(9) That's what she said, too. (a possible reply to Tarek's utterance of (8))

On the one hand, (4)-(5) and (7)-(8) are of the same pattern, but on the other, the report in (6) is much more easily judged correct than is the report in (9). At the end of this chapter, when I lay down my solutions to the problems, I will propose an explanation of this asymmetry. In the meantime, let me explain why these cases pose a problem for the dominant view. The view holds that the proposition expressed is what is said by a given utterance. But (1) and (2) express different propositions, because different times are picked out by the tense. And, however overlapping intervals we take to be picked out by the present progressive in (4) and (5), those, too, end up expressing different propositions. The view is therefore committed to holding that what is said in (1) and (2) is different, and that (4) and (5) say different things, too. This, however, is unsatisfactory, because both our intuitions and our practice of reporting what is said show that there is an important sense in which (1) and (2) say the same thing, and so do (4) and (5).

2.4.3. Is What Is Said Time-Neutral?

At this point, two manoeuvres are available to the defenders of the dominant view. One is to abandon the idea that semantic contents are *eternal* propositions, that is, propositions whose truth only depends on what the *world* is like, and to take them instead to be *temporal* propositions, that is, propositions whose truth may also vary with *times*.[52]

---

[52] Temporal propositions are discussed in A. Prior's pioneering work in tense logic (1957).

I have no objections to temporal propositions *per se* (in fact, we shall see that in my own account, semantic contents receive their truth value as a function of time, too, among other things). What I want to note here is that construing *what is said* as something that may be true at some times and false at others is already a considerable departure from the mainstream view. Recall this well-known passage from Frege:

> But are there not thoughts which are true today but false in six months'
> time? The thought, for example, that the tree there is covered with green
> leaves, will surely be false in six months' time. No, for it is not the same
> thought at all. The words 'This tree is covered with green leaves' are not
> sufficient by themselves to constitute the expression of thought, for the
> time of utterance is involved as well.(1917: 343)

Even Kaplan, who, in his formal system, uses contents that are functions of world-time pairs, gives us, in his informal exposition, every reason to think that the contents expressed by natural language sentences are time-specific.

Switching to temporal propositions may work well enough to deal with the problematic cases that concern the contribution of tense to what is said, but the manoeuvre cannot be easily generalized to other problematic cases, as we will see soon.

The second manoeuvre that may help salvage the mainstream view is to bite the bullet and insist that because (1) and (2) express different propositions, what is said in them is different, and therefore, the report in (3) is literally false, while speakers' intuitions on what is said, according to which (1) and (2) say the same thing, are unreliable and mistaken. As this manoeuvre, unlike the previous one, may be easily replicated in many of the cases that are yet to come, I shall postpone its discussion until section 2.6. We will see, however, that this "literalist" manoeuvre is unsatisfactory.

2.4.4. Is What Is Said Location-Neutral As Well?

Consider Jones in London, on 15 July 2005, saying:

(10) It is cold and windy.

Suppose now that three days later, Miles says in San Francisco:

(11) It is cold and windy.

There is clearly a sense in which Jones and Miles are saying the same thing. For, they are both saying that it is cold and windy. While saying the same thing, namely, that it is cold and windy, Jones is talking of London on the 15th of July, and Miles is talking of San Francisco on the 18th of July, which is why their utterances need not have the same truth value. Still, this difference in truth value does not mean that what is said is different; or at least, there is no obvious reason why it should. Note, though, that when we say the same thing in the sense in which (10) and (11) say the same thing, but are also talking of the same thing, our utterances will have the same truth value.

The mainstream view assigns different contents to (10) and (11). (10) is taken to express the proposition that it is cold and windy in London on July 15, 2005, and (11), the proposition that it is cold and windy in San Francisco on July 18, 2005. What might help here are time-neutral and location-neutral contents, that is, contents that are functions not only of possible worlds, but also of times and places. Now, to account for the intuition that Jones and Miles are saying the same thing, namely that it is cold and windy, a defender of the mainstream view who wants to insist that the propositions expressed are eternal propositions might point out that Jones and Miles are using the same sentence, so that we might report that the same thing has been said simply because the same sentence has been used. Though this account might work fine for (10)-(11), it fails to generalize. Suppose that on Wednesday 20 July, Jones, Miles and Dorsky are together in San Francisco, and it is incredibly cold and windy. The next day in Stanford, Jones says to Miles:

(12) It is cold and windy, though less than yesterday in the city.

Next, suppose that on Saturday July 23, in London, Dorsky says to Miles:

(13) It is cold and windy, though less than Wednesday in San Francisco.

Miles might reply to Dorsky, relying on his earlier conversation with Jones:

(14) Jones said that, too.

The intuitions here are that Miles' reply is just fine, he is not saying anything false. There is certainly a sense in which Jones and Dorsky are saying the same thing, for they are both saying that it is cold and windy, but less than in San Francisco on July 20, 2005. And this might account for the intuition that (14) is true, when properly disambiguated.

Once we accept that there is a sense in which what is said by (12) is the same as what is said by (13), and once we start looking for a semantics of speech reports in which Miles' report in (14) may come out true, we can no longer stay within the mainstream view. For, the propositions expressed by (12) and (13) are different, the one being that Stanford on July 21 is less cold and windy than San Francisco on July 20, the other being that London on July 23 is less cold and windy than San Francisco on July 20. However, unlike (10)-(11), the sentences used by Jones and Dorsky have different lexical meanings. In (12), 'yesterday in the city' bids you to identify the relevant city and the day before the day of utterance. In (13), 'Wednesday in San Francisco' bids you to identify the relevant Wednesday and the place that the name 'San Francisco' stands for. If, for instance, Dorsky were to use the same sentence as Jones, her utterance would be true if on July 23, in London, it was less cold and windy than on *July 22*, presumably again *in London*.

2.5. *De Se* Assertion

The present section addresses the issue of how indexical pronouns contribute to *what is said*. I will point out some interesting asymmetries between the 1st person pronoun and the other pronouns.

2.5.1. How the 1st Person Contributes to What Is Said

In my next challenge to the mainstream view, I will turn against it a case that had earlier served to motivate it. Suppose that Prof. Feferman says:

(1) I am writing a book on Montague.

Next, suppose that, possibly at a different time, Inma says:

(2) I am writing a book on Montague.

Pace Frege's claim that Inma and Feferman express different *thoughts*, and Kaplan's echoing claim that different people using the 1st person pronoun *say different things*, there is an equally strong intuition that, in some important sense, what Feferman said in (1) is the same as what Inma said in (2). Each said that he or she was writing a book on Montague. Thus, having heard Prof. Feferman in (1), I may comment as follows on Inma's utterance of (2):

(3) That's what Professor Feferman said, too.

To be sure, as it stands, (3) is ambiguous between reporting Feferman as having said that Inma was writing a book on Montague, and that he himself was. This ambiguity, in linguists' jargon, is the ambiguity between *strict* vs. *sloppy* readings.[53] Thus, suppose that Prof. Feferman has never heard of Inma, and that this is common knowledge in the context of my utterance of (3). Then the dominant reading is the sloppy reading, on which what Feferman said is that *he*, not Inma, was writing a book on Montague. The problem, of course, is that the semantic contents that the mainstream view assigns to (1) and (2) are different, one involving Feferman and the other Inma, hence either what is

---

[53] E.g. Lasnik (1989), Lappin (1997), Büring (2003). The sloppy/strict distinction is primarily discussed in the literature on ellipsis and anaphora.

said is not the semantic content, or else, one must insist that, notwithstanding intuitions, what is said is different.

Three moves are available to the defenders of the mainstream view. First, one might say that semantic contents are not necessarily *eternal* propositions. In the same way in which temporal propositions were used to deal with those cases in which the time relevant to the truth value of the utterance appeared not to be part of what is said, what we might call *egocentric* propositions can help us deal with those cases in which the *person* relevant to the truth value of the utterance does not appear to be part of what is said.[54] Egocentric propositions are propositions that may take different truth values relative to different individuals; in other words, they are *properties*.

The suggestion that what is said is, at least sometimes, an egocentric proposition was made, for instance, by R. Feldman:

> We can say that what I assert by uttering 'I was insulted yesterday' is something that can be true for, or relative to, one person at one time, while being false relative to some other person at the same or some other time. So you and I assert the same thing by uttering 'I was insulted yesterday' and this thing may be true for me when I assert it and false for you when you assert it. (1980: 79)

What lends further plausibility to the idea that the content of an assertion is an egocentric proposition, or a property, is that the idea squares very well with the view according to which contents of beliefs and other "propositional" attitudes are, in fact, properties. This was the view that David Lewis put forward in order to deal with the problem of *de se* attitudes.[55] What sometimes goes unnoticed in Lewis' account of *de se* attitudes is that there are two crucial components to his view. One is that the content of the attitude is a property (rather than a proposition). The other is that to believe such a content is to *self-*

---

[54] Unsurprisingly, *egocentric* propositions were also pioneered by Prior (1977).

[55] See Lewis (1981). To my knowledge, Lewis did not argue that contents of *assertions* are also properties. Nor did argue that there were not (unless we interpret the earlier quoted passages from Lewis (1980) as claiming that there is no semantic theory of what is said).

ascribe it, that is, to believe that the property applies *to oneself*. This will turn out to be relevant later on, as we shall see that not only in *de se* attitudes, but also in *de se* assertion, this mechanism of *self-ascription* plays a crucial role. In (1), it is not just that Feferman asserts a certain property, viz. the property of writing a book on Montague. He is asserting this property *of himself*.

Now, whether or not egocentric propositions are the right way to handle *de se* assertion and account for the intuition that Feferman in (1) and Inma in (2) are saying the same thing, what should be emphasized is that this manoeuvre really amounts to giving up the mainstream view. The indexical 'I' has been, for Kaplan and his followers, a directly referential expression *par excellence*. It has been seen as the paradigm of expressions that contribute their reference, and nothing but their reference, to the semantic content and to what is said. To say that in (1), Feferman, who is the speaker and therefore the semantic contributions of the word 'I', is not part of the semantic content amounts to giving up one of the central tenets of the received wisdom on indexicality and direct reference.

The second manoeuvre to deal with *de se* assertion, which has become something of a standard response made by those who want to account for the sense in which (1) and (2) say the same thing without stepping out of the Kaplanian tradition, is to point out that the sentences used by Feferman and Inma are the same, and then suggest that when two people use the same sentence or utter the same words, they may be truly reported as having said the same thing. I am going show that this manoeuvre is unsatisfactory for two reasons. First, it fails to generalize. For, an equally strong intuition of same-saying may be triggered, in *de se* assertion, even when the sentences used are neither the same nor synonymous. Second, it over-generates. For, there are cases that minimally differ from (1)-(2) in which speakers use the same words, yet one cannot truly report them as having said the same thing (except by adding special qualifications).

Finally, the third, "literalist" manoeuvre consists again in insisting that the report in (3) is, strictly speaking, false, and that the intuition of same-saying is mistaken. Again, this manoeuvre will only be discussed in section 2.6.

2.5.2. Different Meanings, Different Propositions, Same Thing Said

We have seen that when different people say "I am writing a book," there is a sense in which they are saying the same thing, for each is saying that he or she is writing a book. The usual way of dealing with such cases is to point out that those people are all using the same sentence, which would then explain why we are inclined to hear them as saying the same thing – for, after all, they are uttering the same words. I want to show that this is not the right response to the problem of *de se* assertion. Although using the same sentence may partly account for the intuition that the same thing has been said, that cannot be the end of the story. For, even when the propositions expressed by the two utterances are different, it is neither *necessary* nor *sufficient* to use the same sentence in order to be intuitively saying the same thing, or for the report that the same thing has been said to come out intuitively true.

Suppose that during Feferman's class on Montague, Inma tells Tarek:

(4) I really like this class.

The following week, Tarek and Maria are talking about the classes that they like or dislike, and Maria says:

(5) I really like Feferman's class on Montague.

Tarek may correctly reply to Maria:

(6) Inma said that, too.

The sentences used by Inma and Maria are obviously different, and so are their lexical meanings. Thus imagine that Tarek and Maria's conversation is taking place during Prof. Lawlor's metaphysics class. Then if Maria were to use the same sentence that Inma used, she would have ended up saying that

she really liked Lawlor's metaphysics class, and Tarek's reply that this was also what Inma said would then be false. For, the report in (6) is true only to the extent that both Inma and Maria were talking of one and the same class.

The problem, to sum up, is that the propositional contents that Kaplanian theories assign to (4) and (5) are different, the first being that Inma really likes the class at stake, and the second, that Maria really likes it. But the lexical meanings, or Kaplanian characters, associated with those sentences are also different, hence the truth of Tarek's report in (6) cannot be attributed to Inma and Maria's having uttered the same sentence, since they have not.

2.5.3. The 1st Person vs. the 2nd Person

Just as using the same sentence is not required for the same thing to be said, it is not enough either. Consider the following (minimal) pair of situations:
(i) *de se* assertion

(7) I am too old for graduate school. (Inma talking to Tarek)

(8) I am too old for graduate school. (Maria talking to Tarek)

(9) That's what Inma said, too. (Tarek's reply to Maria)

(ii) *de te* assertion

(10) You are too old for graduate school.
    (Prof. Feferman talking to Inma, overheard by Tarek)

(11) You are too old for graduate school. (Tarek talking to Maria)

(12) (?) That's what Prof. Feferman said, too. (still Tarek talking to Maria)

There is a striking asymmetry between the 1st person and the 2nd person pronoun in how they behave in speech reports. Consider (9). As it stands, it has two readings: one on which Inma is reported as having said that *Maria* is too old for graduate school (the *strict* reading), and one on which she is reported as having said that *she herself* is too old for graduate school (the

sloppy reading). If it is, say, common knowledge in the context of (9) that Inma would have never said such a thing about Maria (say, because she has no idea who Maria is, or because such a comment would have been politically incorrect and Inma would never say such things), then the dominant reading of (9) is its sloppy reading, and (9) comes out true in virtue of Inma's having uttered (7). However, if we try the same sort of sloppy report by simply replacing 'I' by 'you', no such report seems to be available. For, there is a very strong intuition that (12) is not ambiguous, but downright false (assuming that Prof. Feferman never said that *Maria* was too old for graduate school).

This asymmetry between *de se* and what one might call *"de te"* assertion raises the following problem. Suppose, following the mainstream view, that propositional contents play the role of what is said. But (7) and (8) have different contents, and still, in some important sense, they say the same thing. In both cases, the speaker is saying of herself that she is too old for graduate school. Furthermore, the report in (9), when properly disambiguated, is uncontroversially true. Now, one might think that this is because the sentences uttered in (7) and (8) are the same. But consider (10) and (11). Here, too, the sentences uttered are the same, but we do not get a sloppy reading for the report (12). That report is not ambiguous, but false. This shows that something was missing in the account that the mainstream view gave us for the *de se* cases in the first place.

Now, one might object that, after all, there is a sense in which speakers making *de te* assertions and asserting the same thing, not of themselves, but of the person to whom they are talking, are saying the same thing. For, aren't both Feferman in (10) and Tarek in (11) saying that their *addressee* was too old for graduate school? Indeed, there appears to be this sense of same-saying for *de te* assertion as well. And, what's more, it is possible to truly report Feferman in (10) as saying the same thing as Tarek in (11), but under certain conditions. The reporter must *explicitly* mention the relevant addressee:

(13) That's what Prof. Feferman said, too, to Inma. (Tarek talking to Maria)

Unlike (12), which our intuitions tell us is not ambiguous but false, the report in (13) has two readings: one on which Feferman said to Inma that Maria was too old for graduate school (the *strict* reading), and one on which he told her that *she* was too old for it (the *sloppy* reading). Similarly, consider:

(14) Prof. Feferman told Inma that she was too old for graduate school.

(15) That's what Tarek told Maria, too.

If I tell you (14) and you reply with (15), your reply is ambiguous between reporting Tarek as telling Maria that *she* was too old for graduate school, and his telling her that *Inma* was too old for it.[56]

The lesson to be drawn is that the strict/sloppy ambiguity exists for reports of *de te* assertions as well, provided that the person to whom the reportee was talking is explicitly mentioned in the report.[57]

2.5.4. The 1ˢᵗ Person vs. the 3ʳᵈ Person

The difference between the 1ˢᵗ person pronoun and the 3ʳᵈ person pronoun in how they behave in speech reports is even more striking. Consider the following minimal pair with respect to (i) and (ii):

(iii) *de re* assertion

(16) She is too old for graduate school. (Prof. Feferman, talking of Inma)

(17) She is too old for graduate school. (Tarek, talking of Maria)

(18) (?) That's what Professor Feferman said, too. (in reply to Tarek)

---

[56] There is even a third reading, on which what Tarek told Maria is that Feferman told Inma that she was too old for graduate school. But let us put momentarily this third reading aside.
[57] In some cases, there may be a strict/sloppy ambiguity in the report, even if the addressee is not specified. Here is one. On their first date, Tarek tells his girlfriend: "You are terrific." She replies: "That's what every boy says on their first date." What she is likely to be saying is not that every boy says that *she* is terrific, but that they all tell the person they are dating "You are terrific." In this sort of case, the sloppy reading is probably available precisely because the strict reading is pragmatically inhibited.

Again, (18) does not seem to be ambiguous. There is a unique, determinate way of understanding the report, namely, that Feferman said that Maria was too old for graduate school.

To be sure, one might object that even with the 3rd person pronoun, there is a sense of same-saying worth being accounted for. After all, both Feferman in (16) and Tarek in (17) are saying that the person to whom they are referring is too old for graduate school. And again, it is possible to truly report Feferman and Tarek as having said the same thing, provided that we make it explicit that the reportee was referring to someone else:

(19) That's what Professor Feferman said, too, about Inma.
   (a possible reply to Tarek's utterance of (17))

Note that with the 2nd person, when the addressee is explicitly mentioned, we have the strict-sloppy ambiguity exactly as with the 1st person, while in the 3rd person case, where we get some sort of sloppy reading when we specify *about whom* what was said was said, we have no *ambiguity*, since the about-clause forestalls the strict reading. Also, with *de te* assertion, it is crucial that the initial sentences should both contain the 2nd person pronoun, while with *de re* assertion, the use of the pronoun is irrelevant. Indeed, the report in (22) is not any worse than the report in (19):

(20) Inma is too old for graduate school. (said by Feferman)

(21) Maria is too old for graduate school. (said by Tarek)

(22) That's what Professor Feferman said, too, about Inma.
   (a possible reply to Tarek's utterance of (21), on the basis of (20))

Let me take stock. We first saw that in *de se* assertion, that is, in assertion that a speaker makes about himself or herself, there is a strong intuition of same-saying. This is a problem for the mainstream view, because, on the assumption that 'I' is an indexical *par excellence* and contributes the speaker to the semantic content, we are committed to saying that *de se* assertions made by

different speakers inevitably have different semantic contents, so the semantic content cannot be the asserted content. In response to this problem, a defender of the mainstream view could say that the reason why we have the intuition that the same thing has been said, and why we can truly report such speakers as having said the same thing, is that they have used the same sentences. But this response will not work. First, as we saw, it is just as easy to have the same intuition of same-saying when the speakers use different, non-synonymous sentences, and to truly report them as having said the same thing. Second, if using the same sentence, or uttering the same words, is supposed to account for certain cases in which we can correctly report that the same thing has been said, then why is it that in some other cases, such as those involving the 2nd and 3rd person pronouns, discussed in the present subsection, the prediction turns out to be wrong? Until it can provide a satisfactory answer to this question, the mainstream view on what is said proves to be extremely limited.

2.6. A Hopeless Enterprise?

In this section, we will try to see whether there are reasons to doubt that the notion of what is said is robust enough to be captured within a theory of semantic content. Paul Ziff (1972) offered examples showing how versatile the locution 'what is said' can be. Here is an example inspired by one of Ziff's. Suppose that Inma and Tarek went to a certain party, to which Maria could not go. Maria now wants to know how the party went. She asks Inma, who says:

(1) Oh, I shouldn't have gone. 'Twas better to stay home and do some needle-work.

Later, Maria asks Tarek, and he tells her:

(2) My goodness! I don't think I've ever been more bored in my whole life.

Maria might well reply:

(3) Yes, Inma said that, too.

Or, talking yet to someone else about the party, Maria might report:

(4) Both Inma and Tarek told me that the party was really boring.

Although the intuitions are not very robust in the case of (3) and (4), it is still true that in everyday life, we often report people as having said the same thing when all that their utterances have in common is a certain implication, and that implication is relevant in the context of the report. Thus, while Inma does not literally say in (1) that the party was boring, what she does say, viz. that instead of going, she should have stayed home to do needle-work, *implies*, given enough contextual background (such as the assumption that needle-work is not a particularly exciting activity), that the party was boring. And similarly for (2).

Is this "implication sense" of what is said, as Ziff calls it, an insuperable obstacle to a semantic approach to what is said? Not for the defenders of the mainstream view who would adopt the literalist response, and who would say that Maria's reply in (3) is literally false, though it may convey something true, and that the report in (4) is likewise literally false, though it is true that, *loosely* speaking, Inma and Tarek said the same thing, because one could *infer* that the party was boring from what they said *strictly* speaking.

For once, the literalist view might be on the right track. For, even if we have the intuition that the reports (3) and (4) are not exactly false and could communicate something true, we also have the intuition that (1) and (2) do not exactly say the same thing. This intuition, too, is reflected in the way in which we report what others have said:

(5) The party was very boring. (said by Tarek)

(6) Inma said that, too.
   (Maria, on the basis of Inma's utterance of (1))

(7) That's not quite true. Inma only said that she shouldn't have gone, and
    that it was better for her to have stayed home to do some needle-work.
    (a possible reply to Maria)

Intuitively, the reply in (7) is true. This, in turn, suggests that when we judged (3) to be true, our judgments of truth were not very robust. Indeed, once we ask ourselves what it is *exactly* that Inma said, and once we realize that, after all, Inma was perhaps only trying to finish a piece that she was knitting, which might have been the reason why she thought that she should have stayed home, and if so, that the party maybe wasn't even boring for Inma, then we are actually willing to retract our earlier judgment that the report was true and to re-evaluate it as being false. According to the literalist, the report had been false from the start, and the instability of our judgments, they would say, precisely shows that our intuitions on what is said cannot be trusted.

For the literalist, the reports in (3), (4) and (6) are inevitably false. But the position that I would like to defend is slightly more complex. I do not hold that the report in (3) is necessarily false. It might well be true. What "retraction" shows is merely that a report that was true in one context may well turn out false in another context (even if the reported utterances are the same). What accounts for this change of truth value is that the standards for what counts as same-saying are different. In the context of (3), the standards were "low": having some relevant implication in common was enough to make the report true. In the context of (7), where the exact wording has been raised to salience and the standards have been shifted to high, the report in (3), as well as the report in (6), are seen as false. Having noted this, for our present purposes we might just as well endorse the simple-minded literalist view.

Somewhat ironically, the plausibility of the literalist response in this sort of case is precisely what shows it to be a bad response to the cases previously discussed. For, there is a striking difference between the implication sense of

what is said and the sense in which, say, different speakers making the same *de se* assertions say the same thing. Consider:

(8) I am writing a book on Montague. (said by Prof. Feferman)

(9) I am writing a book on Montague. (said by Inma)

(10) Prof. Feferman said that, too.
   (Tarek's reply to Inma, on the basis of Feferman's utterance of (8))

(11) (?) That's not quite true. He only said that *he* was writing a book on
   Montague. (Maria's reply to Tarek)

While we were inclined to judge that, in (7), the "that's not quite true" reply was correct, making in turn the report in (6) false, the case is very different with (11), where we would judge Maria's reply incorrect. More precisely, if someone remarks to us that Tarek's report in (10) is false, our reaction is that the person did not properly disambiguate the report, and that they take it to be false only because they have got the wrong reading. The same thing happens with incorrect disambiguations of the strict/sloppy ambiguity of the more usual type, as with VP-ellipsis:

(12) Inma loves her boyfriend.

(13) So does Maria.

(14) (?) That's not quite true. She only loves her own boyfriend.

When presented with a reply as in (14) to the pair (12)-(13), we will, of course, say that one who made such a reply did not resolve the ellipsis in (14) correctly. Exactly the same explanation goes for Maria's reply in (11). In none among the cases with which I have challenged the mainstream view would we re-evaluate the report as being false, the way we do with the report in (6). In all those cases, once the report is perceived as true, it stays true, while any

"that's not quite true" reply may be discarded as arising from an incorrect disambiguation of the report.

As previously noted, this difference between *loose* and *literal* reports of what is said strongly suggests that those who have dismissed the notion of what is said on the grounds of its being too context-dependent or vague, as was done by Lewis, Salmon, Soames, etc., have been too quick in doing so. The examples discussed above vindicate the distinction between loose and literal uses of the locution 'what is said'. Somewhat ironically, though, what we have just seen is bad news for the literalist response to the cases brought up in 2.4. and 2.5, since those cases precisely fall on the side of *literally true* reports, while the literalist held that they were *literally false* but *conveyed* something true.

2.7. What Is Said as Lexical Meaning

I would now like to present my own theory, which takes what is said to be lexical meaning, and identifies the latter with the semantic content. It will be helpful to start by reminding the reader of the main tenets of my proposal regarding lexical meaning and semantic content, and, at the same time, explain how these notions are related to what is said.

a) What is said by a given utterance is nothing more or less than the lexical meaning (or the semantic content) of the sentence uttered.

b) The lexical meaning of a sentence is something that can be true with respect to some things (in some worlds, at some times) while being false with respect to some other things (in the same or some other world, at the same or some other time).

c) The semantic content of both indexical and non-indexical expressions is their lexical meaning, and semantics works as explained in Chapter 1. In particular, the truth clauses associated with indexicals require double indexing, and when the indexical is in the scope of some operator (including indirect discourse operators), it resets the index of evaluation back to the designated index.

d) When we make an assertion, we do not merely assert the semantic content of the sentence, i.e. the content associated with the sentence in virtue of its lexical meaning. We assert that content *of*, or *about*, something. That of which we assert the content will typically be that to which we are referring.

e) The conditions lexically associated with indexicals, while being fully part of the semantic content, and therefore of the asserted content, are usually only asserted for heuristic purposes, to help the hearer figure out what the speaker is referring to and talking about.

f) When we report what is said by some utterance as being the same (or different) as what is said by some other utterance, we presuppose it known in the context of the report what or whom the reported utterances were respectively about. By default, we presuppose that the reported utterances were about the same thing.

To get a better understanding of my proposal, let us see how it works on an example. Suppose that Inma, holding up a hammer in her hand, says to Tarek:

(1) This is heavy.

In uttering (1), Inma is referring to the hammer: she is drawing Tarek's attention to it, so that he may identify the hammer as that of which she wants to say something. Since the hammer at stake also satisfies the conditions encoded in the lexical meaning of 'this', viz. being something salient and proximal, Inma will manage to say something *about* that very hammer, namely, that it is heavy. But the hammer itself is not part of the semantic content of (1), nor is it part of what Inma said.

Now, the lexical meaning of the whole sentence is something that will be true with respect to some things at some times and false with respect to the same or different things at the same or some other time. For example, it is true of the hammer in Inma's hand and of the situation in which she is uttering (1), but it is false of the same hammer in a situation in which Sonia utters (1) while holding up a feather in her hand (because the hammer is no longer salient, hence it is no longer "a this"), and it is obviously false of the feather itself.

Let us briefly see what tenet (f) amounts to. The suggestion is that a person reporting two utterances as saying the same thing will normally presuppose that the reported utterances are also *about* the same thing, or will at least presuppose that the audience knows what or whom the reported utterances were respectively about. Reconsider Inma's utterance, made in reference to the hammer, followed by a reply from Tarek:

(2) This is heavy.

(3) That's what Maria said, too.

When Tarek reports Maria as having said the same thing as what Inma has just said, the issue of what it was that Maria was talking about will be taken for granted in the context of the report. Given that Tarek does not explicitly mention anything as being the thing of which Maria said that it was heavy, it will be understood that Maria said it of the same thing as Inma did, namely, of that same hammer. This is why the report will come out intuitively false when Tarek makes it on the grounds of Maria's uttering the same sentence as Inma did, while referring to something other than the hammer, as, for example, to her suitcase.

The presupposition that Maria was talking of the same thing as Inma can, of course, be explicitly canceled. That is what we saw in 2.5.4:

(4) This is heavy. (Inma, talking of the hammer)

(5) That's what Maria said, too, of her suitcase.
    (Tarek's reply to Inma)

If Maria did say that her suitcase was heavy, Tarek's reply is, of course, true.

2.8. Solving the Problems

My proposal identifies what is said by a given utterance with the lexical meaning of the sentence uttered. Let me now quickly review the key cases

from previous sections, the cases that motivated the mainstream view as well as those that were obstacles for it, and show how my account handles them.

2.8.1. Different Meanings, Same Things Said

A major challenge to my proposal is that we often intuitively say the same thing using sentences that, lexically, do not mean the same thing:

(1) I have graduated in math. (said by Inma)

(2) She has graduated in math. (said by Tarek pointing at Inma)

How are we going to capture the intuition that, in some important sense, what is said by those utterances is the same, if the meaning of (1) is, roughly, that the *speaker* has graduated in math, while the meaning of (2) is that *some contextually salient female* has. If lexical meaning = what is said, isn't that a straightforward denial that Inma and Tarek have said the same thing?

My proposal, taken at face value, denies indeed that (1) and (2) say the same thing *tout court*. But that does not mean that there is no sense in which Inma and Tarek may be understood and correctly reported as having said the same thing. I account for this in three steps. First, when we look at the sentences uttered in (1) and (2), we see that they contain the same verb phrase, viz. 'has graduated in math'. The sentences uttered do, then, overlap in their lexical meaning (and therefore in what is said) to the extent that they *predicate* the same thing. Secondly, we need to explain how those parts that have different lexical meanings may be ignored in judging whether the same thing has been said. The parts that are ignored correspond to the words 'I' and 'she'. As noted earlier (1.2), the conditions lexically associated with indexicals are asserted primarily for heuristic purposes, to help the hearer figure out what the speaker is talking about.[58] Finally, the third step in our account of why (1)

---

[58] We have also noted, in 1.7, that indexicals could turn out to be presupposition-triggers, and that the conditions lexically encoded in indexicals could be considered *presupposed*. The fact that a given thing satisfies such a condition, like the fact that Inma is the speaker of (1), or that she is the most salient female in the context of (2), are not the sort of facts that we normally assert or communicate, nor do we need to inform other people of such facts. Rather, the speaker will assume that such facts are already known to her audience, and will exploit this

and (2) are so easily heard as saying the same thing lies in the fact that the person that Inma is talking about is the same as the person that Tarek is talking about: it is Inma herself. So, both Inma and Tarek are talking about the same person (Inma), and are asserting about her one and the same thing, namely, that she has graduated in math.

True enough, our account also predicts that there is a sense in which Inma and Tarek are *not* saying the same thing. If you wish, Inma is, after all, saying of herself that she is a speaker (an "I"), while Tarek is saying of her that she is a salient female (a "she"). But, as already pointed out, this difference in what is said is not significant, because Inma is saying of herself that she is a speaker only so that her audience may figure out that she is talking about herself.

2.8.2. The Contribution of Time

It is easy to see that my account of the basic case from 2.1.1. – a case that had served to motivate the mainstream view – applies straightforwardly to the first case from 2.4.1. – a case that casted doubt on that view –, a variant of which is repeated below:

(3) I have graduated in math. (Inma, on March 31, 2006)

(4) Inma has graduated in math. (Tarek, talking of Inma, on April 3, 2006).

(5) She said that, too. (a possible reply to Tarek's utterance of (4))

Although the time is relevant to the truth of (3) and (4), and although it is possible that their truth values diverge (for instance, if Inma graduated only on April 2, (4) is true but (3) is false), it has been noted that, intuitively, (3) and (4) say the same thing, and that the report in (5) is true.

The account that I have provided in the case of (1) and (2) can be replicated to the letter in the case of (3) and (4). The only thing that is worth adding is

---

knowledge to help them figure out about whom she is talking. Given that those parts of the lexical meaning of a sentence that correspond to indexicals are only asserted for such heuristic purposes, it should come as no surprise that they should often be ignored in reporting what is said.

that since (3) and (4) are not directly about time, nor are they about events whose time is a subject matter of the conversation, the report in (5) does not come with the presupposition that when Inma said what Tarek said, she was talking of the same time; for, neither Tarek nor Inma were really talking about the time prior to which Inma is said to have graduated. That is what, I think, accounts for the difference with the dinner case:

(6) I am having dinner. (said by Inma, on March 31, at 7 pm)

(7) Inma is having dinner. (said by Tarek, on March 31, at 11 pm)

(8) She said that, too. (a possible reply to Tarek's utterance of (7))

The report in (8) is easily perceived as false. We perceive it as true, though, if the time is explicitly mentioned:

(9) That's what she said, too, earlier this evening/at 7 pm.
     (possible replies to (7))

To be sure, (9) might well be ambiguous between reporting Inma as having said at 7 pm that Inma was having dinner then, or as having said at 7 pm that Inma would be having dinner around 11 pm.

The explanation of why we are inclined to say that the report in (8) is false is that, since having dinner is an event the time of which seems to be relevant in the context of an utterance such as (7), the report that someone else said the same thing *presupposes* that the other person was also talking of the same time. That is, remember, the gist of the tenet (f) of my theory. Now, given that Inma did not say *about* late evening time, eleven-ish pm, that she was having dinner *then*, we take the report to be false. But, the falsity of the report does not come, then, from some difference in what was actually said, but rather, from the presupposition that it was said about the same thing time. This explanation is further supported by the fact that the presupposition may be explicitly canceled, as in (9).

2.8.3. De Se *Assertion on the Strict Reading*

Another major challenge to my proposal is to account for the fact that often, even though the same sentence has been uttered, there is the intuition that different things have been said, while reporting that the same thing has been said seems incorrect. Let me start with the case from 2.2.2. (originally Kaplan's example, which he took to motivate his view). Consider:

(10) I was insulted yesterday. (said by Kaplan, on Friday, March 31, 2006)

(11) I was insulted yesterday. (said by Maria on Saturday, April 1st, 2006)

(12) Kaplan said that, too. (a possible reply to Maria)

Because (10) may be true while (11) is false and vice versa, Kaplan concludes that what is said must be different: what is said in (10) is about Kaplan and about what happened on March 30, while what is said in (11) is about Maria and about the following day.

In my account, what is said in (10) and in (11) is the same thing: it is the lexical meaning of the sentence "I was insulted yesterday" – roughly, that the speaker was insulted on the day before the day of utterance. However, this meaning is not true or false simpliciter, but it is true or false with respect to a person at a time (and at a world). So (10) can only bear a truth value after being given a person and a time to be evaluated at, and that will presumably be the person and the time that the speaker is talking about and referring to. Kaplan is talking about what happened to him on Thursday, March 30 (he was insulted), while Maria is talking about what happened to her on Friday. The difference in their respective subject matters accounts not only for a possible difference in truth values, but also for the intuition that different things have been said.

So, just to repeat, I submit that what is said in (10) and in (11) is the same, but because it is said of different individuals and of different days, people may get the intuition that what is said is different.

The difference in subject matter is also what accounts for the fact that the report in (12) may be false. Recall that report is ambiguous between a strict and a sloppy reading, but if we put the sloppy reading aside for the moment, the report is false (on the assumption that Kaplan didn't say that Maria was insulted on April 1st). The reason why it is false is that when we report that the same thing has been said, we typically take it for granted that this was also said about the same thing. If it was not, then we ought to make it explicit that it was said about something or someone else. But if we leave that implicit, the truth value of the report will be determined depending on whether the two reportees indeed said the same thing *about one and the same person or thing*.

2.8.4. De Se *Assertion on the Sloppy Reading*

Let us now turn to the sloppy readings of reports of *de se* assertions, that is, reports of utterances of sentences that contain the 1st person pronoun. Let me use my example from 2.5.1, slightly simpler than Kaplan's, since it does not involve temporal indexical such as 'yesterday':

(13) I am writing a book on Montague. (said by Prof. Feferman)

(14) I am writing a book on Montague. (said by Inma)

(15) Prof. Feferman said that, too.
     (Tarek's reply to Inma, on the basis of Feferman's utterance of (13))

As noted repeatedly, (15) is ambiguous. Suppose that it is common knowledge between Inma and Tarek that Prof. Feferman does not know of Inma, hence that he could not have been saying anything specifically about her. Then (15) will be naturally understood as reporting Feferman as having said that *he* was writing a book on Montague. Tenet (f) of our theory predicts that when we report what is said by Inma in (14) as being the same as what Prof. Feferman said, we presuppose it known in the context of the report what Feferman's reported utterance was about. The presupposition that his utterance was about the same person as Inma's, namely, about Inma herself, is pragmatically

canceled, since it is common knowledge that Feferman could not have been talking about Inma, of whom he had never even heard. Given that his utterance was not about Inma, the next most plausible hypothesis was is that his utterance must have been about himself.

More generally, the data on *de se* assertions and our practices of reporting them (but also our practices of reporting *de se* beliefs, for that matter) suggests the following interpretation rule:

g) when, with respect to a given *de se* assertion u_1, we report what is said by some other utterance u_2 as being the same as what is said by u_1, either u_2 is about the same person as u_1, or it is a *de se* assertion, hence it is about the utterer of u_2.

The generalization in (g), together with tenet (f), makes it also possible to account for the same-saying of *de se* assertions with sentences that do not have the same lexical meaning. Consider:

(16) I really like Feferman's class on Montague. (said by Inma)

(17) I really like this class. (said by Maria about Feferman's class on Montague)

(18) Inma said that, too. (Tarek's reply to Maria, on the basis of (16))

In (17), we have two subject matters: Maria, and the class she says she likes, namely Feferman's class on Montague. But (17) is also a *de se* assertion, as it is about the speaker herself, namely Maria. The report in (18) is, of course, ambiguous, but if we put aside the strict reading, the sloppy reading is accounted for by means of rule (g): the utterance made by Inma and reported as saying the same thing as Maria's was also a *de se* assertion. Moreover, the two *de se* assertions assert the same thing, namely, that the speaker really likes a certain class. Tenet (f) further tells us that the report in (18) presupposes that the two assertions, though about different "selves", should otherwise be about the same subject matter, hence about the same class, viz. Feferman's class on Montague.

2.8.5. Some Conclusions

Let us overview some of the most interesting moments of this chapter. One of the central tenets of the mainstream view is the distinction between lexical meaning and what is said. In 2.2, we first saw two types of cases that, at a first glance, motivate this distinction – the "different meanings, same thing said" cases and the "same meanings, different things said" cases. But then, in 2.4. and 2.5, I presented a series of cases that fit very badly into the Kaplanian picture. For, we often get the intuition that the same thing has been said, even though the reported utterances do not have the same Kaplanian content, nor do the sentences uttered have the same lexical meaning. We have also seen that in some cases, namely, in *de se* assertions, utterances of sentences that lexically mean the same thing are easily reported as saying the same thing, even when their Kaplanian contents are different. However, in some other, minimally different cases, namely, in *de te* and *de re* assertions, such reports only work either if the assertions are about the same thing (hence their Kaplanian contents, or at least their Kaplanian *temporal propositions*, are the same), or if the reporter makes it explicit that the reportees were talking about different things. In 2.7, I put forward my own account of what is said, on which what is said is nothing less or more than the lexical meaning of the sentence uttered, which, in the light of Chapter 1, is the semantic content associated with the sentence. The two key notions for my account are lexical meanings (or semantic contents) and that about which the speaker is talking (or that to which the speaker is referring, or the subject matter). Finally, I tried to illustrate how my account works by applying it both to the examples that seem to motivate the mainstream view, and those that appear to undermine it.

————

end of Chapter 2

————

Beyond Indexicals

In this chapter, I show how the view laid down in the previous chapters may be extended to certain non-indexical expressions, and how it sheds light on a number of issues, some of which have been known for a long time, while others have been brought up only recently.

3.1. Definite Descriptions

Definite descriptions are one of the most important topics in the 20th century philosophy of language. The discussion that follows is thus bound to remain somewhat rough. What I would like to do is simply outline the directions that an analysis of definite descriptions would take within the framework laid down in the previous chapters. I should stress that the proposed analysis is just *a* possible analysis, but, as we also will see in the case of proper names, there may be several accounts that are compatible with the account that I am offering for indexicals. For instance, both Russell's original proposal (Russell (1905)) and the generalized quantifiers analysis (e.g. Barwise and Cooper (1981), Neale (1990)) are compatible with the account that I defend for indexicals. However, I think that some of the ideas from Chapters 1 and 2 point in the direction of a somewhat different analysis – one that shares certain features with Russellianism (thus, the description itself will be part of the semantic content), but also makes it possible to account for certain motivations

behind a referentialist approach to definite descriptions, such as the intuition that a speaker using a definite description may not only convey, but *say* something *about* some object or person to which or whom she is referring.

Following the lines of my proposal that semantic content is lexical meaning, I will hold that the semantic content of a definite description is also simply its lexical meaning. In turn, the lexical meaning of a definite description is just the lexical meaning of the description itself. But at this point already, one might object that it is not clear that there is such a thing as the lexical meaning of the description. What is the lexical meaning of a word like 'pencil'? We saw that it approximately means something like "an object with a solid strip that leaves marks used in writing or drawing". But then, it is relatively easy to think of some object that should not count as a pencil yet satisfies the paraphrase, and, the other way round, some pencils may fail to satisfy it. However, this objection need not worry us, since, as I have stressed on several occasions, such descriptions are merely approximations to the lexical meaning. For our purposes, we may simply think of the meaning of a word such as 'pencil' as a function that, given a world and a time, returns some set of objects – intuitively, all of and only the things that count as pencils. So the worry that it is not clear that descriptions have lexical meanings, albeit not spurious *per se*, is spurious in the context of the present discussion.

What is more interesting for us here is that the conditions of existence and uniqueness will not be part of the semantic content of the definite description. This will prove to be a welcome result. My analysis is meant to incorporate the best of Russellianism, while being able to account for certain phenomena that are problematic on Russell's account. Those are, of course, the referential uses of definite descriptions, and the problem of incomplete descriptions. In my account, the attributive use and the referential use are on a par, in the sense that neither is derived from the other. The distinction between the two uses will not amount to some difference in the semantic content. The semantic content of a definite description, whichever way it happens to be used, is

always going to be the same. In this respect, my account is not vulnerable to the objection "from ambiguity," typically made to those who hold, as I do, that whether the description is used attributively or referentially may have an impact on the truth value. Beside the attributive and the referential use, I will also consider the anaphoric use of definite descriptions as well. There, too, what is nice about my account is that anaphorically used definite descriptions have the same semantic content as they do on any other use.[59]

How are we going to account for all those different uses while keeping the semantic content the same? In a nutshell, the differences will be accounted for by manipulating the parameter of the assignment of values to the free variables. So here is the plan for the present section. In the next subsection I will clarify the three uses of definite descriptions, and in 3.1.2, I will show, albeit in rough lines, how these uses get accounted for within my account. In the final subsection (3.1.3), I will briefly consider definite descriptions in what is said, and show that my account scores a point there, too.

3.1.1. Attributive, Referential (Deictic), and Anaphoric Uses

The referential/attributive distinction is, of course, due to Keith Donnellan (1966). I will only recall it briefly, and mention three other distinctions with which it should not be confused.

Consider the following sentence:

(1) Smith's murderer must be insane.

Suppose that inspector Barashvili utters (1) after coming upon Smith's savagely mutilated body. She does not have any idea yet who Smith's murderer might be, and what she is saying – or, if you prefer, what she is trying to communicate to her audience – is that *whoever* murdered Smith must be insane. This is a typical instance of an *attributive* use of the definite

---

[59] Let it be noted, *en passant,* that I am not covering the whole range of phenomena that have been discussed in the case of definite descriptions. For instance, the *functional* uses of definite descriptions, as in "The Department Chair changes every 4 years", will be completely ignored.

description 'Smith's murderer' (or 'the murderer of Smith' – I will treat the two as equivalent, since nothing important hinges on this).

Now suppose that someone, say Olaf, has been charged with Smith's murder, and suppose that during the trial, Olaf behaves foolishly, says stupid things, and the like. Now Barashvili wants to communicate to her interlocutor that this person, Olaf, must be insane. Though Olaf is salient, suppose that he is not salient enough to be immediately picked out as the reference, had she simply used 'he'. This means that she needs to find some other way of making it clear that it is of that man out there, Olaf, that she is saying that he is insane. If, in this situation, she utters the same sentence as in (1), she is intuitively saying *of* Olaf that he is insane. What makes this a typical case of a *referential* use is that there is a specific individual to whom the speaker is referring and of whom she is saying something, and the description is only used for heuristic purposes, to help the hearer figure out who this individual is. As Donnellan points out, Ms. Barashvili might not even know whether Olaf really murdered Smith. But insofar as the description is doing its work of helping the hearer figure out that she is talking about Olaf, it does not matter that much whether he actually satisfies the description or not.

We have just touched upon an issue easily confused with referentiality, namely *misdescription* (or the idea that a referentially used description may fail to apply to the thing successfully referred to). Criticisms of Donnellan have often focused on the cases of misdescription (in particular, Kripke (1979)). It goes without saying, though, that a referentially used description may well apply to the object referred to – and it most cases, it does indeed. Furthermore, we are not committed to saying that an utterance containing a referentially used description that misdescribes the thing referred to should be true, even when this thing satisfies whatever is asserted about it). So suppose that Tarek actually isn't Smith's murderer, but is insane. Then the fact that Barashvili uses the description referentially when she utters (1) at the trial, and manages to inform her interlocutor that Olaf is insane, does not commit us to saying that

(1) is true simpliciter. (More on this in the following subsection.) Finally, there could be misdescription even in the case of attributive use. Thus suppose that Smith was not murdered, but died, say, of heart attack, and once he was already dead, some insane maniac savagely mutilated his body. When Barashvili utters (1) over Smith's dead body, what she really means to be saying is that *the person who did that, whoever he or she might be*, must be insane. The lack of some specific intended referent is a hallmark of an attributive use, but the description 'Smith's murderer' is not satisfied by anyone. So we seem to have a combination of misdescription and attributive use.[60]

Also not to be confused with the referential/attributive distinction is the distinction between definite descriptions that satisfy the uniqueness constraint and those that do not, also called *incomplete* descriptions. Suppose that a UPS delivery guy comes into Joy's office to deliver a package, and she tells him:

(2) You may leave the package on the table.

There are many packages in the universe, and many tables as well. So Russellian accounts, taken at face value, seem to predict that (2) is false, which goes against people's intuitions. Probably the best way of handling incomplete descriptions, and of providing a semantics that gives correct truth prediction without having to give up Russell's analysis, is to make use of the contextual domain restriction (see e.g. Stanley and Szabo (2000)). Contextual restriction on the domain of quantification seems to be needed for other quantifiers as well, and straightforwardly applies to the case of definite descriptions, the idea being that in the context of (2), the restricted domain includes only one table and only one package.

Though the contextual restriction strategy does not work that badly, it has its problems. For instance, when the UPS arrives with a package, there may already be other packages lying around, and there may be more than one table in Joy's office. In such a situation, (2) is still correct. The problem is to explain

---

[60] Of course, one can simply say that in this case, the description is used attributively, but since nothing satisfies it, (1) is false – and I was certainly not saying that (1) should be seen as true.

how the domain of quantification gets to include the package in the deliverer's hands, and the table on which he may leave it, without including those other packages or tables as well. However, here is not a place to criticize (let alone defend) the contextual restriction strategy. What I want to note is that incompleteness is not to be confused with referentiality. First of all, not all referentially used definite descriptions are incomplete. The case of (1) uttered at the trial is a case of a referential use in which the uniqueness condition obtains (if Olaf is the murderer). Secondly, not all incomplete descriptions are used referentially. As an example of an attributively used incomplete description, suppose that we have had a delicious meal in a restaurant, though we do not know who the chef is. Still, you might say:

(3) The chef, whoever he or she might be, is very talented.

So there, the definite description 'the chef' is incomplete (since there are many chefs in the universe), but attributive.[61]

The third distinction not to be confused with the referential/attributive distinction is the wide/narrow scope distinction. It is widely held that definite descriptions are scope-taking devices, which I will take for granted in my discussion. The issue of whether the definite description on a particular occurrence takes wide or narrow scope with respect to some operator is orthogonal to the issue of whether the description is used attributively or referentially.[62]

Let me close this subsection by throwing in yet another phenomenon to be dealt with: the anaphoric uses of definite descriptions. Consider:

(4) If a truck hits a car, the car will be more damaged than the truck.

The definite descriptions 'the car' and 'the truck' are used here as devices of anaphora. The reasons for using a description are clear: if you try to use just a regular pronoun, that would presumably be the pronoun 'it' both times, and

---

[61] A similar case is discussed in Neale (1990).
[62] See e.g. Kripke (1979).

we get something hopelessly ambiguous. If we spoke the language of FOL, we could have dispensed with anaphoric uses of definite descriptions:

(5) If x is a truck and y is a car and x hits y, then y will be more damaged than x.

We will shortly see how our account takes care of the anaphoric uses.

3.1.2. Different Uses, Same Content

My proposal, in a nutshell, is that the semantic contribution of a definite description such as 'the car' is, simply, the semantic content of 'car', which, in turn, is a (suitably constrained) function from world-time pairs to sets of individuals (namely, cars). Let me now apply this proposal to the three uses of definite descriptions that we saw in the previous subsection, beginning with the attributive use. Suppose that inspector Barashvili, staring over Smith's savagely mutilated body, says:

(6) Smith's murderer is insane.

The truth conditions that we want to get for (6), taking ψ for 'murderer of Smith' and ϕ for 'insane', are as follows:

$S, w_1, t_1, w_2, t_2, f \models (\text{the } x: \psi x) \, \phi x$ iff

(i) There is one and only one assignment f' like f except at most for the value for x such that: $S, w_1, t_1, w_2, t_2, f' \models \psi x$.

(ii) $S, w_1, t_1, w_2, t_2, f' \models \phi x$. [63]

Let us now turn to the referential (or deictic) use. Barashvili now wants to say *of Olaf* that he is insane, and the definite description at hand that will help her refer to Olaf is 'Smith's murderer'. So Barashvili utters the same sentence

---

[63] Two remarks are in order. First, we are using the notation (the x: ψx) ϕx, rather than ϕ(the ψ), so that the definite description, *qua* a generalized quantifier, may take different scopes (an issue that matters mainly when we look at the interaction of definite descriptions with modal and other operators). Second, the assumptions of existence and uniqueness only apply to ψ, and so, to avoid any possible misinterpretation, I chose this cumbersome truth clause that clearly marks the separation. Note that the formulation of the truth clause involves something like the "dynamic binding": the assignment function f' in (b) is "bound" by a meta-quantifier over assignment functions from (a).

as in (6). Taking again ψ for 'murderer of Smith' and φ for 'insane', the truth conditions that we now want for (6) are as follows:

S, w_1, t_1, w_2, t_2, f |— (the x: ψx) φx   iff
S, w_1, t_1, w_2, t_2, f |— φx and S, w_1, t_1, w_2, t_2, f |— ψx.

Barashvili's reference to Olaf will, of course, be taken care of through the assignment function f. The assignment relevant to the truth of (6), when the definite description is used referentially, is that which assigns Olaf to x. Note that (the x: ψx) φx, analyzed as above, leaves the variable x free.

Our account of referential uses has an advantage over the standard referentialist accounts with respect to the problem of misdescription. There are basically two ways of running a referentialist account of referential uses. One is the way Kaplan does it, namely, by using the 'dthat' device, which keeps the description at the level of character, so that what gets contributed to the content is that thing or individual that actually satisfies the description.[64] The other gives priority to speaker's reference: it is the thing or individual referred to by the speaker is actually referring that goes into the content.

Now, both of those accounts have a problem when dealing with cases of misdescription. So suppose that Olaf actually isn't the murderer of Smith. Suppose that he is just the guy who savagely mutilated the body, but that he only did it after Smith had been already murdered by Knut. On Kaplan's account, what Barashvili says in (6), using the description referentially, is *that Knut is insane.* But this result is not only counter-intuitive, but also unfaithful to Donnellan's own understanding of the referential/attributive distinction (as noted by Kaplan himself). In the other referentialist account, given that Barashvili is referring to Olaf, (6) is true iff *Olaf is insane*. The problem there is that, since Olaf is not the murderer, we are somewhat reluctant to say that (6) is simply true.

---

[64] What gets contributed to the content is, according to Kaplan, not the thing itself, but its individual concept. But, as noted earlier, we can keep talking the way referentialists do, since nothing important hinges on this.

Donnellan himself, when he talked of misdescription, did not say that an utterance containing a referentially used *mis*description was true simpliciter (assuming that the referent satisfied what is asserted of him/her/it). He only said that by an utterance such as (6), Barashvili may say *something* true. This does not mean that what she said is true as a whole, or that the utterance is true as a whole, but that at least part of what she said is true. Now, my account can easily account for the idea that only a part, but not the whole, of (6) is true. Let me remind you of the truth clause for (6), with 'Smith's murderer' used in reference to Olaf:

S, w_1, t_1, w_2, t_2, f(x)=Olaf |— is-insane(x) and
S, w_1, t_1, w_2, t_2, f(x)=Olaf |— is-a-murderer-of-Smith(x).

The top clause is true, because Olaf is insane – and that is the *something true* that Barashvili said in uttering (6). But the bottom clause is false, so we can account for the intuition that (6), as a whole, is not true. So it appears that, when it comes to misdescription, we can have our cake and eat it, too.

Finally, let us look at the anaphoric use:

(7) If a truck hits a car, the car will be more damaged than the truck.

Ignoring the future tense in the consequent, we want to get the following analysis for (7):

S, w_1, t_1, w_2, t_2, f |— $\forall x \, \forall y((\text{truck}(x) \wedge \text{car}(y) \wedge \text{hits}(x,y)) \rightarrow$
$$(\text{the } z\text{: car } z)(\text{the } v\text{: truck } v)(\text{is-more-damaged-than}(z,v))$$

$$\text{iff}$$

for every assignment f' that is like f except at most for x and y,
if S ... f' |— truck(x) and S ... f' |— car(y) and S ... f' |— hits(x,y),
then S... f' |— more-damaged(z,v) and S... f' |— car(z) and S... f' |— truck(v)

As we could see in 1.8.2, a proper analysis of anaphoric uses requires a mechanism of *unification*. What unification does is to bind variables z and v

that have been left free. The choice is, basically, between unifying x with z and y with v, or unifying x with v and y with z. Here is where the description itself becomes relevant. For, if you unify x with z and y with v, you get (7) false (or more precisely, it is true either if the antecedent fails, hence if no truck ever hits a car, or if trucks are also cars and vice versa, both of which are almost trivially false). That is why we unify x with v and y with z, and that way, (7) ends up being equivalent to the following:

$$\forall x \forall y ((truck(x) \wedge car(y) \wedge hits(x,y)) \rightarrow (is\text{-}more\text{-}damaged\text{-}than(y,x))$$

Though the proposed analysis leaves a good number of issues open, the take-home message should be clear enough: whether the description is used attributively, referentially, or anaphorically, what it contributes to the semantic content does not change. It always contributes the description itself, and that is all that there is to the semantic content. But there are other things that language users need to know in order to be competent with definite descriptions. They need to know what to do with the assignment function. And that is where the distinctions among the different uses show up. On the referential use, definite descriptions work very much like indexicals. One of the differences, at least so far, is that we did not force the interpretation of the description into the designated index, the way we did with indexicals. It would be interesting to see what double indexing may bring to the semantics of definite descriptions. But given that definite descriptions are scope-taking devices, whenever you need to interpret the description back at the designated index, rather than at the current index of evaluation, there is always the option of giving the definite description wide scope over the index-shifting device. On the anaphoric use, definite descriptions work very much like those indexicals that have anaphoric uses, namely, 3rd person pronouns (again, with the difference that for anaphoric pronouns, we used double indexing, which we did not do for anaphorically used definite descriptions). It is on the attributive use that definite descriptions are unlike indexicals, and more like

quantifiers. But this use, too, is to be handled by manipulating the parameter of the assignment function.

In sum, there is no doubt that there are differences among the various uses of definite descriptions, differences that get reflected at the level of truth value. Yet, that does not make definite descriptions ambiguous – their semantic content always remains the same. Taking our claim from 1.8. one step further, we may say that we are making progress towards a unified account of attributive, deictic (or referential), and anaphoric uses.

3.1.3. *Definite Descriptions and* What Is Said

The question of how definite descriptions contribute to what is said may be a book topic of its own. The goal of this subsection is not to study this topic in depth, or to bring out problems about what is said specifically due to definite descriptions, but rather, to propose an account of pairs of sentences such that the one contains a definite description and the other an indexical, and that are easily judged and reported as saying the same thing.

Suppose that we are at a party and a man comes to us and says "Hello, I just wanted to say 'hi'". We had met the man before, but we both forgot his name. After he has left, I tell you:

(8) The man who just came by to say 'hi' is a philosopher of science.

Later you talk with other people, and you are saying to Inma how there aren't that many philosophers at the party. She replies, "well well, there is John and Ken and Graciela and Michael," and then she turns and points to that same man – the one who came to say 'hi' –, and says:

(9) And he, too, is a philosopher of science.

To Inma's utterance of (9), you would truly reply:

(10) That's what Isidora said, too.

How do we account for the fact that your reply in (10) is true, and that, at least in a sense, Inma and I said the same thing, given that the sentence I uttered is different from the sentence that she uttered, and that their lexical meanings are different, too, and that therefore, according to my account, what I said is different from what she said?

We would account for it in the very same way in which we accounted for the same-saying between Inma's utterance of 'I have graduated in math', and Tarek's utterance, in reference to Inma, of 'she has graduated in math' (cf. 2.8.1). Here is, roughly, how the account would go. My sentence in (8) and Inma's sentence in (9) overlap in their lexical meaning – and they do so in an important way, since they overlap on the verb-phrase (i.e. 'is a philosopher of science'). So in this sense, there is something that both Inma and I said: it is the propositional function that takes an individual, a time, and a world as inputs, and returns value True if this individual is a philosopher of science at that world-time pair, and False otherwise. Secondly, that on which the sentences in (8) and (9) diverge in their lexical meanings are the contributions of the definite description 'man who just came by to say 'hi'' and of the pronoun 'he'. It may be noted, though, that the latter contributes the condition of being male, which already figures among the conditions contributed by the definite description. So then, I did say in (8) all that was said in (9), but I said something more: I also said of that man that he came by to say 'hi'. To be sure, I only said this of him to help you figure out of whom I was talking. At any rate, the point is that the diverging parts of what is said in (8) and (9) are only asserted for heuristic purposes, and because of this, they tend to be ignored in judging about and in reporting what is said. Finally, Inma and I were both talking of the same guy, which is why you can truly reply to her "That's what Isidora said, too", without having to specify of whom I said it. You simply presuppose that I said it of the same person as Inma did. This final point is, again, an application of tenet (f) of our theory (cf. 2.7).

It is when the description is used referentially that it is easiest to get true same-saying reports. But even when the description is used attributively, we may have the intuition and truly report that the same thing has been said. Suppose thus that Prof. Feferman shows me some graduate student's paper on Montague that he has blind-read, and tells me the following:

(11) The person who wrote this paper, whoever he or she might be, is very smart.

I recognize that it is Inma's paper. Later I see Inma, and I tell her:

(12) You are very smart. In fact, that's what Feferman said, too.

My report that Feferman said the same thing as I did in (12) is intuitively true, even though the definite description in (11) was not used referentially.

We can reuse some of the account given for the previous case. For, the sentence that Feferman used in (11) and mine in (12) overlap in their lexical meaning in exactly the same way as those in (8) and (9) did. They overlap on the verb-phrase, and that is important because that is the "core" of the assertion. As to the parts on which our sentences diverge in meaning, we know why the conditions lexically encoded in 'you' may be ignored – they were only asserted of Inma in order to help Inma realize that I was talking of her. But in (11), the lexical meaning in the definite description, because it is used attributively, does not have the same heuristic character. Since Feferman isn't directly referring to Inma (he doesn't know that she wrote the paper), he can't be using the definite description to help me figure out that he is referring to Inma. *En passant*, this is exactly the reason why it is much easier to get intuitions of same-saying when the descriptions are used referentially than when they are used attributively.

Here is a plausible explanation of why the condition encoded in the definite description in (11) may, after all, be ignored. The point is that we, who are judging that the same thing has been said, know who wrote that paper on

Montague, so we know that the definite description uniquely applies to Inma. From our point of view, the content of the description is no longer crucial, given that we know who is the unique person to whom this content applies. That is why we stop paying attention to that part of the lexical meaning of the sentence in (11). Note, though, that if Feferman overheard my utterance of (12), he could object to my report, protest that he never said that *Inma* was smart, or at least ask "When did I possibly say that?"

To account for the true reading of the report in (12), we can reapply the account given for the previous case, insofar as the presupposition of the same subject matter is concerned. Recall that what tenet (f) of our theory said is that in reporting two utterances as saying the same thing, the reporter presupposes that the reported utterances are about the same thing. When they are not, one will explicitly mention their respective subject matters. Now, Feferman's utterance of (11) and mine of (12) are about the same person, namely Inma. To be sure, we acknowledged that Feferman was not referring to Inma. So the sense in which he was talking *about* Inma is perhaps not exactly the same as the sense in which I was talking about her when I was addressing her directly. Feferman's utterance is about Inma to the extent that she uniquely satisfies the description that he used, so it is *de facto* about Inma. And here again, we can see that the reporter's knowledge, on the one hand, and the knowledge of the persons who are judging whether the same thing has been said, on the other, play important roles. I, *qua* reporter in (12), could presuppose that Feferman's utterance of (11) was about the same person as mine only because I knew that Inma was the person who wrote that paper on Montague. And the intuition that my report in (12) is true is partly due to our knowledge that Inma is the only person who satisfies the description used by Feferman.

3.2. Predicates of Taste

While definite descriptions are one of the most discussed topics in philosophy of language even, predicates of personal taste (words such as 'delicious', 'fun', etc.) have received a fair amount of attention in the past few years. What I want to do in this subsection is show that such predicates may be nicely incorporated into my framework.

3.2.1. Preliminary Remarks

For convenience, I will treat predicates of taste as coming with an implicit argument: a word like 'delicious' will be treated not as a one-place but as a two-place predicate, and its second argument will receive as its value the "taste value" relevant to the truth of the sentence containing the predicate. However, it has been argued, e.g. by Kölbel (2002) and Lasersohn (2005), that the correct semantics for taste predicates *must* use contents that are functions of a taste parameter, which sits side by side with the parameters of world and time, and that the taste parameter *cannot* be seen as an implicit argument lexically associated with the predicate. I am going to show that their claim is wrong. The two *prima facie* rival approaches – the *implicit argument* approach, which takes the taste parameter to be an implicit argument to the taste predicate, and the *"relativist"* approach, which takes it to be a parameter of the index of evaluation – are, from the viewpoint of semantics, notational variants of one another. More precisely, given any sentence containing a taste predicate, and given any context, the two semantic accounts will predict the same truth value, and are, in that sense, semantically equivalent. The exposition of the two formal frameworks and the proof of their equivalence are relegated to Appendix 2.[65]

The section is organized as follows. I start by presenting an interesting puzzle to which the predicates of taste give rise. The puzzle involves a

---

[65] The "relativist" semantics considered in Appendix 2 is that of Lasersohn and Kölbel, which is admittedly different from the relativist account of John MacFarlane. The latter is immune to the argument from semantic equivalence presented in Appendix 2.

phenomenon that came to be called "faultless disagreement." I will show, at the very outset, that the relativist account of faultless disagreement fails to provide a solution to the puzzle. But that the relativist cannot account for faultless disagreement is is not necessarily bad news for the relativist, because, as I will argue in 3.2.3., there is no such thing as faultless disagreement to begin with. More precisely, I will argue that either the two parties genuinely disagree, hence if one is right, then the other is wrong, or the two parties are both right, but their "disagreement" boils down to a misunderstanding. Finally, in 3.2.4, we will see how predicates of taste contribute to what is said.

3.2.2. The Puzzle of Faultless Disagreement

Consider this dialogue between Tarek and Inma, who have just tasted some soybean ice-cream:

(1) This is delicious.

(2) No, that's not true. This isn't delicious at all.

On the one hand, we are inclined to say that Tarek and Inma disagree. After all, Tarek says that the soybean ice-cream is delicious, and Inma says precisely the opposite. But on the other hand, we are also inclined to say that Tarek and Inma may both be right, and that their seemingly contradictory utterances may be true together. So here comes a puzzle:

a: For any two utterances u_1 and u_2, the utterer of u_1 disagrees with the utterer of u_2 if, and only if, the one denies what the other asserts, which, in turn, means that if u_1 is true, then u_2 is false, and if u_1 is false, then u_2 is true.

b: The utterer of 1 (Tarek) disagrees with the utterer of 2 (Inma).

c: On the assumption that Tarek finds the soybean ice-cream delicious, and that Inma does not, 1 is a true utterance, and so is 2.

The problem is that a, b and c, while independently plausible, lead to a contradiction. For the relativists like Lasersohn or Kölbel, the way out of the puzzle lies in rejecting a. Lasersohn thus writes:

> If you say that roller coasters are fun, and I say they are not, I am negating the same content which you assert, and directly contradicting you. Nonetheless, both our utterances can be true (relative to their separate contexts)." (2005: 645)

My own way out of the puzzle will be to reject either b or c. I suggest that the puzzle stems from an equivocation regarding the term 'disagreement'. Suppose that Tarek utters 1 while holding a cherry in his fingers, even though he means to be saying of the ice-cream that it is delicious, and suppose that when Inma denies that it is delicious, she means to be denying that the cherry is delicious, thinking that it is what Tarek was referring to. Inma and Tarek disagree in the sense that neither takes the other to be saying something true. However, this "disagreement" is nothing more than a misunderstanding. When the two parties genuinely disagree, we want it to be the case that if one party is right, then the other is wrong. So I accept that a dialogue such as Inma's and Tarek's triggers both the intuition that if what Tarek says is true, then what Inma says can't be true, and the intuition that Tarek and Inma, if sincere, are both saying true things. But the two intuitions go obviously against each other, and to see which intuition wins, so to speak, one needs further information on the context in which the dialogue is taking place. If the two parties mean to be expressing judgments whose truth depends on their own taste, then they may both be right, but their disagreement is spurious. If, on the other hand, they intend their judgments to apply universally, to people in general, then they disagree indeed, but only one party will get it right.

Let me call "taste-neutral" those contents that are functions of a taste parameter. Now, one thing that I would like to emphasize at the very outset is that the relativist "solution" to the puzzle in terms of taste-neutral contents is

not much of a solution at all. For, assume that Tarek's utterance of 1 expresses a taste-neutral content, and suppose that, in the actual world, it takes value True when evaluated at Tarek, and False when evaluated at Inma. If Tarek is a competent speaker and is aware that different people like different things, he must be aware that the taste-neutral content that he is asserting is true or false depending on the value assigned to the taste parameter. In the parallel modal case, we are aware that the content that we assert may be true in the actual world while false at some counterfactual state of affairs. We often forget about this, because the world that we are talking about is, by default, the actual world. Nevertheless, it sometimes happens that we intend our assertions to be evaluated for their truth at some other world. This happens in counterfactual reasoning, as when we start by saying "suppose that...." and we continue in the indicative mood. Then all the assertions that follow will be evaluated at the (closest) world in which the supposition holds, and not at the actual world itself. Similarly, when we talk about fiction, as when we say that Sherlock Holmes lived in Baker Street, the asserted sentence is to be evaluated for truth at the world of fiction, rather than at our own reality.[66]

Now, if we assume that speakers are not "semantically blind" about taste predicates, then in uttering 1, Tarek must intend his claim to be evaluated not just at some particular world (presumably the actual world), but also at some particular taste. If Tarek intends his claim to be evaluated at his own taste, and if Inma intends her claim to be evaluated at her own taste, then the relativist has given us no account of their *disagreement*. (Of course, if Tarek intends his claim to be evaluated at some universal standards of taste, and if Inma does so, too, then they genuinely disagree, but only one of them will be right.)

3.2.3. Disagreement vs. Misunderstanding

In the previous subsection, I argued that even if we were to take the phenomenon of "faultless disagreement" for granted, taste-neutral contents would not account for it, because a competent speaker who asserts such a

---

[66] For an analysis of fictional discourse along these lines, see e.g. Predelli (2005: 54).

taste-neutral content must intend this content to be evaluated for truth with respect to some particular tast. Hence if Tarek asserts a taste-neutral content that he intends to be evaluated at his own taste, and if Inma asserts the negation of that content while intending it to be evaluated at her own taste, then we shouldn't say that they disagree. They will only disagree if the one's assertion and the other's denial of one and the same taste-neutral content are also to be evaluated at one and the same taste. But then only one of the two parties will be saying something true.

Now, it should be unsurprising that relativist semantics cannot account for faultless disagreement, because there is no faultless disagreement, as I am now going to argue. Consider the following dialogue:

(3) Paris is lovely. (said by Tarek)

(4) Oh, come on, that's not true. Paris isn't lovely at all. (said by Inma)

Recall that the puzzle was that the two equally plausible intuitions that arise in relation to such dialogues lead to contradiction. One is that Tarek and Inma disagree. After all, Tarek says that Paris is lovely, and Inma says the opposite, namely, that Paris is not lovely. And if they disagree, then if one is right, the other must be wrong. But this conflicts with the other intuition, namely that if Tarek and Inma are sincere, and if Tarek finds Paris lovely and Inma does not, then they are both right and are both saying something true.

I will try to show that a more careful analysis of the data demonstrates that, given any case of *prima facie* disagreement, further considerations about the context and the ways in which the conversation might develop make it possible to decide whether it is a case of genuine disagreement, or only disagreement that arises from a misunderstanding.

Let us first note that the fact of acknowledging that one has been talking about one's own taste shows that there was no genuine disagreement in the first place. Thus one way for Inma and Tarek to resolve their disagreement

would be to recognize that what may be lovely for the one need not be lovely for the other:

(5) OK, all that I was saying is that Paris is lovely for me. (said by Tarek)

(6) OK, and what I was saying is that for me, Paris isn't lovely at all.
(said by Inma)

In (5) and (6), Tarek and Inma have made it explicit that in claiming that Paris is lovely, they were talking of themselves and of their own taste. In other words, in (5) and (6) they have articulated the value for the taste parameter. This they can do on the assumption that the loveliness of something depends not only on the thing itself, but also on one's taste. By acknowledging that in (3) and (4) they were talking about their respective taste, Inma and Tarek will realize that, even though apparently saying contradictory things, their claims may be simultaneously true. And if both may be true, then they are not contradicting each other – as they realize themselves in (5) and (6), having moved from "Oh yes/Oh no" dialogue to "OK/OK" dialogue.

Now, even though, in matters of taste, people sometimes reach some kind of agreement by realizing that they like different things and that their taste partly determines the truth of their claims, at other times they continue disagreeing, as if there were a matter of fact as to whether a given thing had a given property (such as loveliness) or not. In Tarek and Inma's episode, it is easy to imagine that they may never resolve their disagreement, even after they have realized that the truth of their utterances *may* vary with taste:

(7) Paris is lovely! And it's not just that *I* find it lovely; Paris is lovely *tout court*.

(8) No, that's not true. Paris is not lovely – though I understand that for many people it may be.

What do people who disagree whether something is lovely *tout court* actually disagree about, if the truth of what they are saying depends indeed

on a taste parameter? For, if Inma and Tarek were to supply different values for this parameter, and if their seemingly contradictory utterances happen to be both true because of being evaluated at different values, there would be no matter for disagreement. A plausible answer is that in the case of genuine disagreement, the two parties agree on the value supplied to the taste parameter, and disagree as to whether, given that value, something is lovely or not. The two parties cannot be both right, and in this respect, their disagreement is on a par with ordinary disagreement about facts. If Inma claims that Maria has never been to Paris, and Tarek says that she has, then they disagree as to whether Maria has ever been to Paris or not. That is a matter of fact, and only one of them can be right, depending on whether Maria has indeed ever been to Paris. Similarly, if Tarek says that Paris is lovely, *simply lovely*, and Inma denies that, what is likely to be at stake, I suggest, is that they intend to be making *generic* claims. Tarek is roughly saying that Paris is lovely for people in general, and that is what Inma is denying (though she accepts that it may be lovely for many people). To resolve this sort of disagreement, they would need to determine whether Paris is lovely on such universal standards, perhaps by some sort of survey.[67] Once this is determined, the one who is wrong should admit that she was wrong and that her claim was false, and accordingly revise her beliefs. For instance, in the example considered, it would be reasonable for Inma to admit that Paris *is* lovely according to some generally accepted standards, and that hers are an exception – just as she would decide, if shown a picture of Maria taken in Paris, that Maria has been to Paris after all, and that her earlier claim was false.

---

[67] To determine whether something is lovely, or tasty, or fun, *in general*, a survey will work only if we assume the 1st person authority, that is, if we assume that everyone knows what is lovely/tasty/fun for himself or herself. Expressions of taste do appear to obey the first person authority in the way in which belief reports do. If I sincerely report my own beliefs, then it seems that I cannot go wrong. Similarly, if we are deciding the issue of whether something is tasty *to me*, it seems plausible to think that no one can tell better than myself whether or not that is the case. However, whether this assumption is correct or not is not a properly *semantic* issue. Both the classical implicit argument semantics and the relativist semantics discussed in Appendix 2 have the advantage of staying neutral on the issue of what it is that determines whether or not something is fun, delicious or lovely with respect to a given taste.

Finally, disagreement, in the sense of taking the other party to be saying something false, sometimes comes from a misunderstanding. Let us turn back to the dialogue between Inma and Tarek, repeated below:

(9) Paris is lovely.

(10) Oh, come on, that's not true. Paris isn't lovely at all.

Our immediate reaction, as previously noted, is to take Tarek and Inma to be disagreeing with each other. But suppose that to Tarek's insisting that Paris is surely lovely, Inma replies:

(11) Paris is all but lovely. She is arrogant and fake.

It is at this point that Tarek realizes that he and Inma are not talking about the same Paris. While he meant to be talking of the French capital, Inma, who had been reading the magazine *Vogue*, took him to be talking of Paris Hilton, whose photo was on the cover page.

It happens often enough that people engage in a debate simply because they think that they are talking of the same thing, while they are actually talking about different things. This is what happens in (9) and (10). Tarek, thinking of where he would like to go, was talking of the city, but in doing so, he overlooked the fact that Inma was reading about Paris Hilton and was therefore likely to take him to be talking of *that Paris*, not of the city.

The lesson to be drawn from the case of (9) and (10) is that it often happens that people disagree, in the sense that they take the other person to be saying something false, when there are virtually no facts that they disagree about. Or, more accurately, the facts that their disagreement turns upon are facts about the context and their use of words. In our example, Tarek believes that his use of the name 'Paris' refers unambiguously to the French capital, while Inma believes that Tarek used the name 'Paris' to talk about Paris Hilton. So even here, there is something factual that Inma and Tarek disagree about, namely,

the issue of which Paris is being referred to with the name 'Paris'.[68] But, of course, they do not realize that their disagreement arises from a lack of agreement on what their words stand for. At any rate, such disagreement that stems from a misunderstanding should certainly not qualify as genuine disagreement. And I submit that disputes about matters of taste often conceal such a misunderstanding.

3.2.4. *Predicates of Taste and* What Is Said

Let me take stock. I have argued that there is no such thing as faultless disagreement, and that even if there were such a phenomenon, the notion of taste-neutral contents would not account for it. Forthcoming in Appendix 2 is a proof that the implicit argument semantics and the relativist semantics are notational variants of one another. Nevertheless, all of this does not mean yet that there might not be other considerations that tell in favor or against the one or the other approach. In this subsection, I want to look at some considerations related to what is said that, at a first glance, appear to tell against the implicit argument approach. However, I will argue that those considerations are far from being conclusive. I will also outline how the theory of what is said proposed in Chapter 2 handles the data, regardless of whether we use implicit arguments or taste-neutral contents.

It is often held that the implicit argument account of taste predicates commits you to a wrong construal of the notion of what is said. Let me explain the problem on an example. Consider:

(12) This is delicious. (said by Tarek referring to the chocolate cake)

(13) This is delicious. (said by Inma referring to the same cake)

(14) This is delicious. (said by Tarek referring to the ice-cream)

---

[68] Or, if you think that the name of the French capital and the name of Miss Hilton are actually different, albeit homophonic names, then Tarek and Inma will disagree on the question of which of the two names occurs in the sentence that Tarek uttered.

Let us suppose that in the three cases, the speakers clearly mean to be expressing their own taste. Now arguably, there is a strong intuition that Tarek in (12) and Inma in (13) have said the same thing, namely, that the cake is delicious, and there is arguably an equally strong intuition that what Tarek said in (12) is different from what he said in (14).

Let us, for the sake of the argument, take those intuitions for granted. Then they might seem to pose a problem for the implicit argument account for the following reason. The semantic representation associated with the three utterances is something like *delicious (x, y)*, while the relevant assignments are:

For (12): f(x)=chocolate cake, f(y)=Tarek;

For (13): f′(x)=chocolate cake, f′(y)=Inma;

For (14): f″(x)=ice-cream, f″(y)=Tarek.

But then, (12) and (14) are as similar as (12) and (13). However, the intuition is that (12) and (13) say the same thing, while (12) and (14) do not. The implicit argument account appears to be unable to account for this difference.

By contrast, a relativist might claim that this is not at all a problem for his or her account, because what is said, on the relativist account, is the taste-neutral content, of which the value of the taste parameter is not part, while the values assigned to the variables are. This sort of content is the same for (12) and (13): it is the singular taste-neutral proposition that the chocolate cake is delicious. On the other hand, the contents of (12) and (14) are different, the content of (14) being the singular taste-neutral proposition that the ice-cream is delicious.

In the light of our discussion from Chapter 2, it should be clear that these considerations do not cut sharply in favor of the relativist account, nor do they undermine the implicit argument account. I have argued that a technical notion of what is said that accounts best both for the intuitive notion of what is said and for the use of the locution 'what is said' is the notion of lexical meaning itself. This commits me to holding that what Tarek said in both (12) and (14) is actually one and the same thing – it is the lexical meaning of the sentence 'this is delicious'. But I may happily embrace this consequence of my

view, since, as I showed in 2.7, I have means of explaining the intuition that Tarek said different things. That intuition is explained by the fact that Tarek was *talking about* different things, and was asserting deliciousness *of* those different things, namely, the cake vs. the ice-cream. So in sum, Tarek said the same thing, but he said it *about* different things.

The fact that Tarek was talking about different things also explains why simply reporting him to have said the same thing twice will not work. The reason is that when we report what is said as being the same, we assume that it was said about the same thing. And indeed, if we cancel this presupposition by making it explicit what Tarek was talking about on each occasion, the report becomes correct:

(15) This is delicious. (said by Tarek, talking about the ice-cream)

(16) You said that, too, about the chocolate-cake.
    (a correct reply to Tarek, on the basis of (12))

If we accept the account of what is said from Chapter 2, motivated by independent considerations, then what remains to be explained is why when we truly report Tarek and Inma as having said the same thing in (12) and (13), when Tarek is expressing his own taste and Inma her own, we do not need to make it explicit whose taste each was talking about. Now, recall that the same thing happens more generally whenever what one is talking about is oneself, that is, in the cases of *de se* assertion. Suppose that Tarek says:

(17) I like soybean ice-cream.

Now imagine that Inma, on some later occasion, utters that same sentence. Then one can truly reply to Inma: "Tarek said that, too." To be sure, the reply, as it stands, is ambiguous between the sloppy and the strict reading, but with enough contextual background, the sloppy reading can be made dominant. In any case, the point is that in reporting Inma to have said the same thing as Tarek did, we do not have to make it explicit that Tarek was talking of himself.

The general rule (our rule (g) from 2.8.4) is that when a given speaker A is talking of himself, and you report another speaker B as having said the same thing, it is presupposed that B was talking either of the same thing as A (hence of A himself), or that she was talking of herself. With this rule in place, we can straightforwardly account for the fact that when Tarek, talking of himself and referring to the chocolate cake, utters the sentence "This is delicious" in (12), and when Inma, talking of herself and of the same cake, utters the same sentence in (13), not only have they said the same thing, but it is also correct to report them as having done so, without having to point out that different values for the taste parameter were relevant to the truth of what they were saying. For, after all, talking about one's own taste is a way of talking about oneself.

One might object that this cannot be the end of the story, because, when Inma utters (13) and someone replies to her, "That's what Tarek said, too," the reply is not ambiguous. The reading on which Tarek would be saying that *according to Inma's taste*, the cake is delicious, does not seem to be available. One possible response is that the unavailability of this reading is a pragmatic matter, since it is very unlikely that Tarek would be making a claim whose truth depends on Inma's taste. Another possible response, further developed in 3.3.8. (for the similar case of epistemic modals), is that in (12), (13) and (14), neither Tarek nor Inma are really talking *about* their taste, at least not in the same sense in which they are talking about the ice-cream or the chocolate cake. If this is son, then in reporting that Tarek in (12) and Inma in (13) said the same thing, the presupposition that they were talking about *the same taste*, required for the strict reading. does not even apply.

3.3. Epistemic Modals

Epistemic modals are similar to indexicals in that the truth value of a sentence containing an epistemic modal depends on the context in a non-trivial way. It

does not only depend on the time and the world, but also on something else: it depends on the epistemic evidence that the speaker has, or that she shares with her audience, or, at any rate, on *someone's* epistemic evidence. Exactly *whose* epistemic evidence is relevant to the truth of a sentence containing an epistemic modal is the point on which there is considerable divergence among the existing theories. The account that I will propose might be superior to any among these to the extent that the question "whose evidence?" becomes independent of the semantics of epistemic modals (even though, of course, different epistemic evidence may give rise to different truth values). Just as the question of what is being referred to by a speaker who is using the demonstrative 'this' is not to be settled by semantics – since the meaning of 'this' merely helps constrain the range of potential referents – the question of whose evidence is relevant to the truth of a sentence containing an epistemic modal is similarly settled at the level of pragmatics. Again, the choice of evidence may affect the truth value, but the point is that the choice itself is not a semantic matter.

I will shortly show in greater detail how my proposal can be made to work, and how it shares certain features both with the "relativist" accounts such as John MacFarlane's and with the more traditional "contextualist" accounts.[69] What it shares with relativism is that the epistemic evidence relevant to the truth value is neither part of the semantic content nor part of what is said, and that the same sentence, even if you hold its context of utterance fixed, may be endowed with different truth values if it is evaluated relative to different epistemic evidence. What it shares with contextualism is that the context-dependency of epistemic modals is handled by the same semantic apparatus set up for other context-sensitive expressions, such as indexicals.[70] But before

---

[69] In the literature on epistemic modals, the term 'contextualism' stands for those accounts that hold that the truth value of a sentence containing an epistemic modal depends on some element supplied by the context of utterance. But this context-dependence may be lexically encoded, so in this respect, contextualism is epistemology has little or nothing to do with contextualism in the semantics/pragmatics debate. See 2.3.2.

[70] To which extent this is really a feature shared with the existing contextualist accounts is debatable, since those are essentially built upon the Kaplanian account of indexicals.

we get to the proposal itself, let us start by laying down the data to be accounted for.

3.3.1. Context-Sensitivity and Epistemic Disagreement

Suppose that we are at the department, and you ask me if I know where Prof. Cheng is. I tell you:

(1) She might be in her office, though she might also be in the library.

I do not really know where Prof. Cheng is, but I know that whenever she is around, she is usually in her office or in the library, and I have no reason to think that cheng might be away. Intuitively, my utterance of (1) is true. But now, suppose that Prof. Cheng is actually at a conference in Lima, and, as she is delivering her talk, Inma, who is in the audience, utters the same sentence as I did in (1) – referring, of course, to Prof. Cheng. Now, Inma knows full well that Prof. Cheng is neither in her office nor in the library, since she is right there delivering her talk. Then Inma's utterance of (1) is intuitively false. What this shows is that epistemic modals are context-sensitive, but in a way that goes beyond sensitivity to the world and to the time. For, the world of both Inma's and my utterance of (1) is just the actual world, and the times are also the same (or so we assume).

Let us call the parameter on which the truth of sentences containing epistemic modals turns upon, the *evidence* parameter. What was shown in the case of predicates of taste (the technical results are in Appendix 2) applies, *mutatis mutandis,* to epistemic modals: it does not really matter whether the evidence parameter is treated as a new parameter in the index of evaluation, side by side with the world, the time, and the assignment function, or is taken care of by means of an implicit argument lexically associated with the modal expression. I will opt for the latter. But let us not anticipate too much, and let us go back to the data to be accounted for.

Again, suppose that we are at the department, you ask me where Prof. Cheng is, and, since my evidence does not exclude her being in her office, which is where she is most often when she is around, I say:

(2) She might be in her office.

But you have actually just checked in her office, and you know that she isn't there. So you say:

(3) No, she must be somewhere else.

The problem is that, insofar as my own epistemic evidence is concerned, nothing has changed between my utterance of (1), which we said was true, and my utterance of (2). However, we are no longer prepared to say that my utterance of (2) is true. For one thing, what you say in (3) appears to be a denial of what I say in (2). For another thing, your utterance in (3) is intuitively true. Hence we ought to conclude that my utterance of (2) is false. And yet, if my utterance of (1) was true, and nothing has changed in the world or in my epistemic evidence, then it is unclear why my utterance of (2) should not be true. Again, we have a puzzle.

Note that there is a difference between epistemic disagreement, as in (2)-(3), and disagreement on matters of taste that we discussed in section 3.2. When Tarek says that the soy-bean ice-cream is delicious, and Inma says that it isn't, we, observing from the outside, are willing to grant to both of them that they are saying something true. That is why we were led to conclude that their disagreement was spurious and relied on a misunderstanding. On the other hand, when I say that Cheng might be in her office, and you say that she must be somewhere else, an outside observer who is able to compare your evidence with mine and to establish that you have better evidence will judge that I am saying something false, and that you are saying something true. The intuition is, then, that we genuinely disagreement here, and that only one of us has gotten it right (viz. you).

Things are more complicated when we look at dialogues where the speaker and the hearer have conflicting evidence, and where the outside observers is unable to establish whose evidence is better. Thus, suppose that you know that Cheng was eager to go to the conference in Lima, while I know that she was thinking of cancelling her trip. Your evidence tells you that she must be in Lima, my evidence tells me that she might still be around.

(4) Prof. Cheng might be in her office.

(5) No, she must be at the conference in Lima.

Unlike the case of (2)-(3), we are no longer in the position of establishing who has better evidence, you or me. Nevertheless – and this is what marks the contrast between epistemic disagreement and disagreement on matters of taste – there is still the intuition is that a dialogue such as (4)-(5) reflects genuine disagreement, and that either (4) is true or (5) is true, but not both.

3.3.2. Third-Party Assessment, Retrospective Assessment, and Retraction

The phenomenon of the "third party assessment", following MacFarlane's terminology, could already be seen in the case of epistemic disagreement. The 3rd party assessment is, simply, the fact that when an outsider wonders whether a sentence containing an epistemic modal is true or false, she will care little about the epistemic evidence that was available to the speaker and her audience. Instead, she will rely on her own evidence. In other words, if you know that Cheng is at a conference in Lima, then you will consider my claim that she might be in her office false, regardless of the fact that *I* don't know that she is in Lima, and that you know that *my* evidence does not exclude the possibility of her being in her office.

Retrospective assessment is just like third party assessment, except that the "third" party is the speaker herself, but she is looking back at her utterance with hindsight. So suppose that on Monday I say:

(6) Prof. Cheng might be in her office.

But on Tuesday, Cheng, who has just come back from the conference, tells me that she was in Lima the day before. Arguably, when I learn that she was in Lima, so that she could not have been in her office after all, I will reassess my utterance of (6) as being false.

Let us look more closely at what exactly happens in such cases of retrospective assessment. Consider:

(7) Prof. Cheng might be in her office. (said by me on Monday)

(8) That's false. (said again by me on Tuesday, as a comment on (7))

Recall that 'that' as used in (8) is a device of sentential anaphora, and I have previously stressed that sentential anaphora is a topic still under study (cf. 1.6.5). Now, when in (8) I say "that's false" with respect to (7), an interesting thing happens. While the value for the time parameter that is relevant to the truth of (8) is, so to speak, inherited from (7),[71] the value for the epistemic evidence parameter is not inherited in this way. For if it were, (8) would come out false, because with respect to the epistemic evidence that was available to me when I uttered (7), (7) is true. So the lesson to be drawn is that the time parameter and the evidence parameter, at least in the case of (7)-(8), behave quite differently.

Let me close this subsection with the phenomenon of retraction. Retraction is similar to retrospective assessment in that it typically involves some speaker who is looking back at some utterance that she had made earlier. At the time when she made it, her utterance seemed to be true, but with hindsight, the speaker considers her utterance as being, or as having been, false. It thus seems that whenever there is retraction, there is retrospective assessment as well. But in retraction, the speaker will also judge that she was *wrong* to say what she had said:

---

[71] If it were not, then the time of (8) itself would be relevant to the truth value of (8). But then, (8) would come out false, because on Tuesday, Cheng *might* be in her office (and we may even suppose that when I say (8), she *is* in her office.

(9) Prof. Cheng might be in her office. (said on Monday)

(10) I was wrong to say that. (said on Tuesday, regarding (9))

While we are inclined to think that (8) is true – that is, that with the hindsight, I can say that my utterance of (7) was false, we are less prepared to say that (10) is also true – for, I was not really wrong to say (9): when Cheng is around, she is most often in her office, and having no idea that she was going to Lima, it seems that I did nothing wrong in claiming that she *might be* in her office.

Retraction is thus a somewhat more subtle phenomenon than retrospective reassessment, since it involves issues having to do with justification and responsibility in making assertions. Also, the differences in the ways of retracting one's earlier claim are difficult to get a solid grip on. Thus consider the following as possible comments on (9):

(11) I wasn't wrong to say that, and so I stand by what I said.

(12) I wasn't wrong to say that, but I don't stand by what I said.

On the one hand, if one judges that one was not wrong in making a claim, then it seems that one ought to stand by the claim one made. But after learning that Cheng was in Lima, so that she could not have been in her office, it looks like I should not stand by the claim that I made on Monday, when I said that she might be in her office. So both (11) and (12) are problematic.

The subtleties that we encounter with retraction, and that we do not find with mere retrospective assessment, are the reason why I will focus on the latter. There might be many other interesting phenomena related to epistemic modals, but what we have so far will keep us busy for the rest of this section.

3.3.3. A Solution in Terms of "Indexed" Modalities

The account of epistemic modals that I am proposing is cast within a very traditional account of modalities in general, and epistemic modalities in particular. For the epistemic modal 'it might be that' we will use the "existential" modal operator $\Diamond$, the same one that we used for 'it might be

that' understood as 'it is *metaphysically* possible that', and with the very same truth clause. For 'it must be that' we will, of course, use the "universal" modal operator $\Box$. The trick is that we are going to *index* those modalities. It is an old trick: doxastic operators and epistemic operators such as 'Tarek believes that' or 'Inma knows that' are precisely such indexed (universal) operators in the framework of epistemic logic.[72]

Turning back to epistemic modals, the truth clause for an indexed diamond $\Diamond_a$ deploys a doubly indexed accessibility relation $R_{(a,t)}$. This will slightly complicate the semantics, because the structure of interpretation will have to supply not just one, but a whole bunch of accessibility relations on possible worlds. But this indexing will prove useful in several respects. First, we will have one and the same truth clause for all epistemic modals 'it might be that', regardless of who utters it and when and what their evidence is. Second, the semantic content of an epistemic modal is going to be the same regardless of what the epistemic evidence relevant to the truth value is. Third, we will index the epistemic modalities with free variables. The parametrization of truth will thus remain exactly as before, because the context-sensitivity of epistemic modals will be taken care of through the parameters of time and of assignment of values to the free variables. Now that I have told you, in a nutshell, what the theory is, let me spell it out and then apply it to the data gathered in the two previous subsections.[73]

The truth clauses for indexed epistemic modalities are as follows:

$S, w_1, t_1, w_2, t_2, f \vDash \Diamond_x \varphi$ iff for some w' s.t. $w_1 R_{(f(x),t1)} w'$: $S, w', t_1, w_2, t_2, f \vDash \varphi$

$S, w_1, t_1, w_2, t_2, f \vDash \Box_y \varphi$ iff for every w' s.t. $w_1 R_{(f(y),t1)} w'$: $S, w', t_1, w_2, t_2, f \vDash \varphi$

This means, as already mentioned, that the structure of interpretation S is such that for every $u \in U$ and every $t \in T$, it specifies a relation $R_{(u,t)} \subseteq W \times W$. This

---

[72] See e.g. Hintikka (19), van Benthem (1997).
[73] In the formal framework in Appendix 1, we will, for the sake of simplicity, forget about epistemic modals. However, the formal semantics of epistemic modals provided in the present subsection may be straightforwardly incorporated into the final framework.

is probably more than we need, because even for some silly object such as an empty jar of peanut butter, there will be a relation that obtains when one possible world is accessible from the other given the epistemic evidence available to that jar! However, this overgeneration of accessibility relations is quite harmless.

The proposed truth clauses account for the context-sensitivity of epistemic modals in an obvious way. When I utter a sentence such as back in (1), and we want our semantic system to predict that the sentence is true, we simply need to give it as input an assignment function f which assigns *me* (i.e. the speaker of (1)) to the variable by which the modality is indexed. And when Inma utters (1) in the context of the conference, and we take her sentence to be false, the relevant assignment of values will be the one that assigns *Inma* to the variable indexing the modal operator.

Turning to epistemic disagreement, in dialogues such as (2)-(3) or (4)-(5), we will index the two modalities with one and the same variable, and we will make sure that the value assigned to this variable does not change from the evaluation of the one sentence to the evaluation of the other.[74]

3.3.4. How the Indexing of Modalities Helps in Handling Reassessment

Let us begin with third party reassessment. What happens there is that the speaker evaluates her sentence with respect to one assignment function (in general, the one that assigns the speaker herself to the variable indexing the modality), while the third party evaluates the same sentence with respect to a different assignment function. So in this sense, you might say that my account is "relativist" – one and the same utterance gets different truth values depending on who is truth-evaluating it. But this form of relativism is very conservative. It does not require any changes to the semantic apparatus that we have had from the beginning. What is more, when I obtain that my

---

[74] Note that we are using variables whose values are individuals, even though there are cases in which the relevant value for the evidence parameter is more plausibly the evidence available to a *group* of individuals (such as the speaker and her audience).

utterance of "Cheng might be in her office" is true because I evaluate it with respect to the assignment of *myself* to the variable indexing the modality, and when Inma obtains that my utterance is false because she evaluates it with respect to the assignment of *herself* to that same variable, the case is not that much different than when two people obtain different truth values when they evaluate some utterance with respect to different assignments simply as a result of confusion or miscommunication. Thus suppose that I say:

(13) She is a mathematician.

You take me to be talking about Maria, because she happens to be visually salient in the context of my utterance. I take myself to be talking about Inma, because we have just been talking about her. The assignment with respect to which you evaluate my sentence for truth assigns Maria to the free variable in the sentence, the assignment with respect to which I evaluate it for truth assigns Inma to that variable. We get different truth values, because Inma is a mathematician but Maria isn't. So we do not need to go as far as epistemic modals to get the sort of relativism that is inherent to my account.[75]

Finally, retrospective assessment deploys the time parameter with which epistemic accessibility relations are indexed. Recall the example of (7)-(8), repeated as (14)-(15):

(14) Prof. Cheng might be in her office. (said by me on Monday 1/8/07)

(15) That's false. (said by me on Tuesday 1/9/07, as a comment on (14))

For both (14) and my reassessment of (14) in (15), the relevant assignment of values is the one that assigns *me* to the variable that indexes the epistemic

---

[75] In the next subsection, though, we will look more carefully at the differences between indexicals and epistemic modals. Also, the case of (13) is similar to the case of "Paris is lovely" where one person is talking about the French capital and the other about Paris Hilton. Two out of three theories of proper names that I will consider in 3.4, namely, the metalinguistic theory (my version) and the "no semantics" theory, would both analyze the Paris case along the same lines: the speaker takes the relevant assignment to be the one that assigns the French capital to the free variable in the sentence, the hearer assumes that the relevant assignment is the one that assigns Miss Hilton to that variable.

modal, so that will not account for the change in truth value of (14). What will account for this change is that different times determine different accessibility relations. While the time relevant to the truth of (14) (as assessed on Monday) is Monday, the time relevant to the reassessment claim in (15) will be Tuesday. While there is a possible world in which Cheng is in her office on Monday, and which is accessible from the actual world via the relation $R_{(Isidora,\ 1/8/2007)}$, no such world is accessible from the actual world via the relation $R_{(Isidora,\ 1/10/2007)}$.

The case of (14)-(15) is also interesting for another reason. I have just said how the time parameter relevant to the truth of the sentence in (14) changes from the context of (14) to the context of (15). But this change in time does not affect the issue of *when it is* that Cheng might be in her office. So suppose that the right way of understanding (15) is, as we said, the one on which (15) comes out true, because I now know that Cheng was in Lima on Monday and that, being in Lima, she could not have been in her office.[76]

Here is where double indexing might help. What we need to say is that the present tense in 'it might be' in (14) is the *indexical* present tense, and that the formal rendering of (14) involves the indexical operator 'now'. So the sentence that is uttered in (14), and evaluated as true from the context of (14), and that is reassessed as false from the context of (15), is:

(α) $\Diamond_x$ now (Cheng is in her office).

The relevant values for truth evaluation are as follows:

In the case of (14): w_1:= the actual world, **t_1:=Monday 1/8/07,**
w_2:= the actual world, t_2:=Monday 1/8/07, f(x)=Isidora;

In the case of (15): w_1:= the actual world, **t_1:=Tuesday 1/9/07,**
w_2:= the actual world, t_2:=Monday 1/8/07, f(x)=Isidora;

---

[76] I think, though, that (15) is ambiguous. There is also a sloppy reading (with respect to the present tense) on which what it roughly says is that it is false on Tuesday that Cheng might be in her office (hence, on Tuesday). And since she is in her office on Tuesday, (15) comes out false. I will leave this reading aside, even though my proposal can nicely account for it (which I leave an exercise to the reader).

So, when reassessing (14) from the context of (15), the designated index remains what it was initially. When we evaluate (α) with respect to the values relevant in (15), it is t_1 (Tuesday) that determines the accessibility relation for the epistemic modality, while it is t_2 (Monday), deployed by the indexical 'now', that determines the time at which Cheng is in her office. Thus (α), evaluated from (15), is true iff there is some world w accessible from the actual world via $R_{(Isidora, Tuesday (=t2))}$ such that the pair (w, Monday ($=t_1$)) sends the two-place predicate 'to be in' to a set that includes the pair (Cheng, Cheng's office).

3.3.5. *The Differences Between Indexicals and Epistemic Modals*

Let us note three differences between indexicals and epistemic modals (though there might be more). First, indexicals deploy the designated index in their truth clause, while epistemic modals do not. We could already see how this difference plays out when we discussed the case of (14)-(15), but we will look at it again, and present some cases where the index relevant to the epistemic modal is the current index of evaluation, determined by the interpretation of some operator in whose scope the epistemic modal lies, rather than the designated index. Second, it appears that 'That's not true' and similar expressions that involve sentential anaphora inherit the assignment of values to the variables from the antecedent utterance when the variable stands for an indexical, but not when the variable indexes the epistemic modality. Third, indexicals and epistemic modals behave differently in reports of what is said. The second and the third difference will be discussed in separate subsections.

As we saw in Chapter 1, and as has been often noted in the literature, indexicals are unaffected by the interpretation of intensional operators in whose scope they lie. That is because what matters to the interpretation of the indexical is the designated index, while intensional operators leave this index untouched – they only shift the index of evaluation. Now, the truth clause that I proposed for epistemic modals does not deploy the designated index, but the index of evaluation. Let me remind you of the truth clause:

$$S, w_1, t_1, w_2, t_2, f \vDash \Diamond_x \varphi \text{ iff for some } w' \text{ s.t. } w_1 R_{(f(x), \mathbf{t1})} w' : S, w', t_1, w_2, t_2, f \vDash \varphi$$

Note that the time and world deployed by the accessibility relation are t_1 rather than t_2, and w_1 rather than w_2. This means that we should expect that when an epistemic modal lies in the scope of some modal or temporal operator, that will affect its interpretation. And it does. Though the data are not as clear-cut as one might wish, it is at least true that the embedding operator *may* affect the interpretation of the embedded modal.

Let us start with belief reports, where what we seem to have is not just that the doxastic operator *may* affect the interpretation of the epistemic modal, but that it normally will. Suppose that Inma, who is with Prof. Cheng in Lima and knows that Cheng cannot be in her office, says the following about Maria:

(16) She thinks that Prof. Cheng might be in her office.

Let us go through the derivation that my account provides for (16). Let w_A stand for the actual world, t for the time of (16), φ for 'Cheng is in her office', and let the relevant assignment f be such that f(x)=f(y)=Maria. The truth clause for the doxastic operator 'believes that' was given in 1.6.3 ('$w \succ_{\text{Maria}} v$' means that v is among Maria's doxastic alternatives for w). So here is what we get:

$S, w_A, t, w_A, t, f \vDash$ believes ($\text{she}_x, (\Diamond_y \varphi)$)   iff

$S, w_A, t, w_A, t, f \vDash$ female(x), and:

$S, w_A, t, w_A, t, f \vDash$ believes ($x, (\Diamond_y \varphi)$)   iff

for every w' s.t. $w_A \succ_{\text{Maria}} w' : S, w', t, w_A, t, f \vDash \Diamond_y \varphi$ iff

for every w' s.t. $w_A \succ_{\text{Maria}} w'$ there is some w'' such that $w' R_{(\text{Maria, } t)} w''$ and:

$S, w'', t, w_A, t, f \vDash \varphi$

The reader is invited to verify that this is indeed what we want. On the other hand, if the truth clause of the epistemic modal deployed the designated world, rather than the world of evaluation, (16) would come out equivalent to

"Cheng might be in her office", because the resetting of the index of evaluation to the designated world would just cancel the interpretation of the doxastic modality 'believes that', which operates on the world of evaluation.

Now, there is one respect in which the proposed analysis of (16) might not yet be exactly what we want. As we laid it out, it looks as if it were a sheer coincidence that the value to the variable that indexes the epistemic modal, alias y, is the same as the value assigned to the variable associated with 'she', alias x. But it seems that we should have more than mere "co-reference", to use an old, though for my framework inadequate term. What we ought to have is at least some form of anaphora, if not outright binding. Let me explain. When Inma utters (16), it seems that she will never be able to get her interlocutor to evaluate (16) with respect to an assignment of values that assigns Maria to x, while assigning Inma herself to y. But this is puzzling, because if Inma simply says "Cheng might be in her office", or uses any unembedded epistemic modal, typically, the relevant assignment of value to the variable indexing the modality (alias y) will be the one that assigns Inma. So it looks as if the syntax of (16) tells you that the two variables should be given the same value – which is what we get if we use one and the same variable from the beginning:[77]

$$\text{believes (she}_x, (\Diamond_x \varphi))$$

3.3.6. Tensed Epistemic Modals

It might be worth noting that in a theory such as in Schlenker (2003), the data from the previous subsections would not inevitably show that epistemic modals are not indexicals. They could be indexicals, but *shiftable* ones (unlike the indexical 'I' in English). Schlenker provides an account in which belief operators are monsters, i.e. operators that shift the context (or the designated index), while modal and temporal operators are just the usual kind of

---

[77] This point is even more obvious when the epistemic modal lies in the scope of a doxastic or epistemic operator used with a quantifier phrase, as in: "Everyone thinks that Prof. Cheng might be in her office."

intensional operators, i.e. those that shift the circumstance (or the index) of evaluation.[78] Now, if it turned out that epistemic modals are also affected by the interpretation of temporal operators, that would show that they not indexicals, not even the Schlenkerian shiftable ones. Let us examine the interaction between epistemic modals and temporal operators, in order to see whether we were right to to deploy the time of evaluation, rather than the designated time, in their truth clause.

Suppose that on Tuesday you ask me why on Monday I said that Prof. Cheng might be in her office, given that she was at a conference in Lima. I could reply that I did not know that she was in Lima, or I could simply say:

(17) Well, she might have been in her office!

Intuitively, (17) is true. This shows that the past tense operator overscopes, at least in this case, the modal operator. The analysis of (17) is as follows:

$S, w_A, \text{Tuesday}, w_A, \text{Tuesday}, f(x)=\text{Isidora} \vDash \text{sometime-past } \Diamond_x\varphi$ iff

for some $t<\text{Tuesday}$: $S, w_A, t, w_A, \text{Tuesday}, f \vDash \Diamond_x\varphi$ iff

for some $t<\text{Tuesday}$ and some w s.t. $w_A R_{(\text{Isidora},t)}w$: $S, w, t, w_A, \text{Tuesday}, f \vDash \varphi$

Of course, if we want the relevant past time t to be that Monday, we will either use contextual restriction on the domain of times over which the temporal operator sometime-past quantifies, or, based on the discussion from 1.8.6, we might prefer the following analysis for (17):

$S, w_A, \text{Monday}, w_A, \text{Tuesday}, f(x)=\text{Isidora} \vDash (\text{past} \wedge \Diamond_x\varphi)$

iff for some w s.t. $w_A R_{(\text{Isidora},\text{Monday})}w$: $S, w, \text{Monday}, w_A, \text{Tuesday}, f \vDash \varphi$

and $S, w_A, \text{Monday}, w_A, \text{Tuesday}, f \vDash \text{past}$

---

[78] Schlenker's account is motivated by a study of indexicals in other languages, such as Russian (where the present tense appears to work like a shiftable indexical) and Amharic (in which the 1st person pronoun, in the context of a belief report, may stand for the person whose belief is being reported, rather than for the speaker).

What matters is not how you would analyze the past tense itself, but the question of which day, Monday or Tuesday, is relevant to the epistemic modal. If it is Tuesday, then (17) is false, because there is no possible world in which Cheng was in her office on Monday that is accessible from the actual world via the relation $R_{(Isidora,Tuesday)}$ (since on Tuesday, I know that she was not in her office on Monday). In order to get a true reading for (17), the relevant time must be Monday.

Von Fintel and Gillies (ms.) discuss some cases that seem to suggest that the sentence in (17) also has a reading on which (17) comes out false. We may happily acknowledge this. All we need is to uncover a scope ambiguity in the expression 'might have been'. In the analysis on which (17) was true, we parsed it into sometime-past$\Diamond_x\varphi$. In the analysis in which it comes out false, we will parse (17) into $\Diamond_x$sometime-past φ.[79]

3.3.7. *Indexicals, Epistemic Modals, and Sentential Anaphora*

Another often noted difference between indexicals and epistemic modals is that when you assess or reassess a sentence containing an indexical, you are supposed to keep the original world, time and assignment function. In other words, you evaluate the conditions lexically encoded in the indexical at the world, time and assignment that were relevant in the context of utterance. On the other hand, when reassessing a sentence containing an epistemic modal, you are allowed to supply a new assignment, a new time, and, presumably, a new world as well. The following example will illustrate the difference:

(18) She is a mathematician. (in reference to Inma)

(19) That's false.

Suppose that Inma is a mathematician. Then (19) is false. The fact that one could evaluate the sentence in (18) with respect to the assignment that assigns

---

[79] Or, if we opt for the non-quantificational analysis of the past tense, the scope ambiguity will be between (past $\land \Diamond_x\varphi$) and $\Diamond_x$(past $\land \varphi$).

Maria to the free variable associated with the pronoun, and that Maria is not a mathematician, will not render (19) true. On the other hand, consider:

(20) Prof. Cheng might be in her office. (said by Maria)

(21) That's false. (said by Inma)

Even if Inma knows that Maria's epistemic evidence did not exclude the possibility of Cheng's being in her office, hence Inma knows that (20) is true from Maria's point of view, she can still consistently say that "that" is false, because she may evaluate the content associated with the sentence in (20) with respect to her own epistemic evidence.

Two points deserve clarification. First, when one re-assesses a sentence containing an epistemic modal, one need not keep the original assignment function, *but one may*. In other words, suppose that you truth-evaluate an epistemic modal sentence made by Inma. You may supply your own epistemic evidence, the way Inma did in (21), but you may also decide to evaluate it with respect to Maria's evidence – that is, for instance, what we *qua* theoreticians do. More generally, any person who is not "semantically blind" about epistemic modals, and who knows that their truth depends on an evidence parameter, will have a choice when deciding which value for this parameter to consider as relevant for the truth value. For instance, suppose that I know that Cheng was not eager to go to Lima and that she probably stayed at home, but I also know that Inma had every reason to think that Cheng would be at the conference, and suppose that Inma says:

(22) Prof. Cheng must be at the conference in Lima.

If I were asked whether what Inma said in (22) is true, my answer will depend on whose epistemic evidence I choose to evaluate it at. If the topic of our conversation is the issue of whether Inma was lying, or trying to mislead us into thinking that Cheng was at the conference while she herself did not think so, then I would probably choose to evaluate it with respect to Inma's

own evidence, and I might say "Inma was saying truth". On the other hand, if the topic of our conversation are Cheng and here whereabouts, then I would evaluate (22) with respect to my own evidence, and tell you that what Inma said was false.

The second point of clarification is that even with indexicals, it may happen that we evaluate the sentence with respect to an assignment that is different from the one originally intended. Suppose that I have been talking about Inma with Tarek, and I tell him the following:

(23) She is a mathematician.

Suppose furthermore that you have overheard our conversation, and that, in the situation in which I utter (23), Maria happens to be the most salient woman – we may even suppose that both Tarek and I are staring at her. Still, Tarek and I agree that the relevant value for 'she' is Inma, while you take Maria to be the relevant value. Here we have a third party assessment of a sentence containing an indexical that does not keep the original assignment. Now, this sort of case will typically happen as a result of misunderstanding (even though you undersood what I said; you just misidentified the person of whom I said it.) But one may also decide to evaluate the sentence with respect to a different assignment. In such a case, one will normally make it explicit that a new assignment has been supplied. So, for example, one may truly comment on (23) as follows:

(24) That's true of Inma, but false of Maria.

3.3.8. *Epistemic Modals and* What Is Said

The evidence from reports of what is said only confirms the points made in the last subsection, namely, that when you assess a sentence containing an indexical, you will typically keep the original assignment of values, but you need not; the latter happens either because of a misunderstanding, or because you will explicitly indicate the newly supplied assignment. On the other hand,

when you assess a sentence containing an epistemic modal, you will typically evaluate it with respect to your current epistemic evidence, thus supplying a new assignment of values to the variables, even though you may also decide to evaluate it with respect to the epistemic evidence that was relevant at the time of the utterance. The only novelty that comes with reports of what is said is that the differences between indexicals and epistemic modals are even less striking, because of reports of *de se* assertions. Compare:

(i) *de re* assertion

> (25) He's a loser. (Inma, talking about Tarek, overheard by me)

> (26) He's a loser. (Maria, talking about Aaron, overheard by me)

> (27) # Inma said that, too. (my reply to Maria, on the basis of (25))

(ii) *de se* assertion

> (28) I'm a loser. (Tarek, talking to me)

> (29) I'm a loser. (Aaron, talking to me)

> (30) Tarek said that, too. (my reply to Aaron, on the basis of (28))

(iii) epistemic modal assertion

> (31) Cheng might get an offer from Harvard. (said by Tarek)

> (32) Cheng might get an offer from Harvard. (said by Maria)

> (33) Tarek said that, too.

The empirical data are as follows. (27) is bad. In fact, as it is, it is false. But, as stressed several times in Chapter 2, its falsity does not lie in there being a difference between what is said in (25) and what is said in (26). Rather, it lies in the fact that in replying (27), I presuppose that Inma was talking about the same person as Maria, and that she said *of him* that which Maria said – and this is false, since she never said that *Aaron* was a loser. This presupposition, as

we also saw, may be cancelled, as when I say "Inma said that, too, about Tarek". Turning now to (31), that report is ambiguous. Either it reports Tarek as having said that he was a loser, or that Aaron was a loser. Finally, (34) is not ambiguous – for, it has no false reading, only a true one, and its truth is not affected by the difference between Tarek's and Maria's respective epistemic evidence.

Here is how we account for the data. What is said in (32) and (33) *is* the same thing. The two sentences lexically mean the same thing, so what is said is the same. What might potentially be a problem is that the presupposition that their speakers were talking of the same thing need not be cancelled – at least, not explicitly. There are two ways of handling this problem. Either as with did in the case of taste predicates, which is to make the interpretation rule (g) from 2.8.4. apply more broadly, including not only *de se* assertion proper, but also claims about (one's own) taste and about (one's own) epistemic evidence. And, as a matter of fact, talking about one's own taste or evidence is a way of talking about oneself. Another way of handling the problem is just to point out that when one makes an epistemic modal claim, one is not really *talking about* one's own epistemic evidence. One would be talking about it if one said, for instance, "*For all I know*, Cheng might be in her office". But if there is no explicit mention of one's knowledge or evidence, there is no reason to suppose that one is talking about his or her evidence, so the presupposition of the same subject-matter (= tenet (f) of our theory) does not even apply.

3.4. Proper Names

If I have postponed saying anything about proper names until now, that was for a good reason. Proper names are such a delicate topic, and now that I can no longer keep postponing it, as this work is reaching towards its end, I am turning to it with some unease.

I will actually abstain from defending any particular view on proper names. Instead, I will discuss three views that are all compatible with my account of indexicals. One is the mainstream referentialist view, inspired by Kripke's work and pursued by Donnellan, Kaplan, Salmon, and many others. I will try to show that it is coherent to endorse a referentialist approach to proper names while giving it up in the case of indexicals. However, I will also say why I do not think that, among the three views, this one fits best with my view on indexicals. The second view to be discussed is the so-called metalinguistic view. That one fits well with some aspects of my own proposal in the case of indexicals, but, on the other hand, it has problems of its own, which is why I am not eager to endorse it either. Finally, the third view to be considered shares some ideas with the referentialist approaches, while being significantly different from them, since it holds that proper names do not properly have any semantic content. I will therefore call it the "no semantics" view. I do not know of anyone who explicitly holds this view.[80] Nevertheless, the "no semantics" view might fit well with my general approach, since I identify semantic content with lexical meaning, and it is plausible to think that proper names do not have lexical meanings. But the view appears to entail that proper names cannot be an object of study for semantics, and it has some other potentially problematic consequences as well. Though I will do my best to give some plausibility to this view, I am not prepared to fully defend it yet.[81] The final thought of this section will, then, be that when it comes to proper names, we are already so much beyond indexicals that we are not committed to any particular view. Before I set out to explain and discuss the three views, I want to look at the empirical data to be accounted for.

---

[80] Some writings of John Stuart Mill may be interpreted as anticipating the "no semantics" view, and so may some writings of Peter Geach. In the more recent literature, Dever (1998) suggests that proper names should be semantically treated as free variables, while Vallée (ms.) flirts with the idea that proper names might not even be part of the language.

[81] In fact, there is yet a fourth, well-known approach to proper names compatible with my approach to indexicals: it is the Frege-Russell approach itself. Thus, if proper names turned out to be definite descriptions in disguise, all that I would need to say is that the semantic content of such a name is the lexical meaning of the description itself. As I find the view implausible (partly for the reasons that Kripke (1980) points out), I will not explicitly discuss this view.

3.4.1. Is There Any Lexical Knowledge Associated With Proper Names?

One of the main questions is whether it makes any sense to talk of *lexical* knowledge when it comes to proper names. The plausibility, for instance, of the metalinguistic view, turns precisely upon one's answer to this question. Related to it is the question of what kind of *logical* inferences one is entitled to make from sentences containing proper names. The third and last issue that I will discuss concerns the contribution of proper names to what is said.

Let us start with the question of whether proper names have lexical meanings. At a first glance, you might think that it is obvious that they do not. For instance, the mere fact that you will not find in a dictionary an entry such as 'Tarek Al'Mouhari' suggests that this is not a word endowed with any lexical meaning – for, after all, that is precisely what dictionaries are there for, to tell you what the words of a given language mean.

Even if the right answer to our question may well be that proper names have no lexical meaning at all, it is not as obvious as you might think. It has been suggested by philosophers such as François Recanati that there is a notion of lexical meaning that you may plausibly associate with proper names. The idea is that the mere fact that some string of symbols or sounds is a *name* endows this name with the meaning 'whoever bears that name'. So, the lexical meaning of any name N would be obtained by means of the following schema:

(the metalinguistic thesis) The lexical meaning of 'N' is "the bearer of 'N'."

Regardless of whether the metalinguistic thesis is correct or not, there is yet a third possible answer to the question of whether proper names have lexical meanings. The answer is that some names do, and some don't. E.g. 'Tarek Al'Mouhari' clearly does not. But it might be suggested that there are names such that a person who ignores their meaning should not really count as a competent speaker of English. For example, the names of the twelve months in the calendar ('January', etc.), the names of the days in the week ('Sunday', etc.), names such as 'Earth', 'Sun', 'Moon', the names of the nine planets, the

names of the continents ('Africa', 'Asia'...), and so on, seem to be good candidates for being names knowledge of whose meaning is constitutive of the knowledge of English. These are proper names, yet a person who does not master them may be plausibly charged of not being a fully competent English speaker. The problem with this view is, of course, where to draw the line. For instance, you might think that 'France' is such a name, that it has a lexical meaning, and that speakers are not fully competent in English if they lack the knowledge that France is a country, that the adjective derived from 'France' is 'French', etc. But, by parity of reasoning, the name of any country should then be part of the English lexicon. Yet, if someone has never heard, say, of Oman, and does not know that the derivative adjective is 'Omani', you may charge this person of being uneducated, but it does not seem quite right to charge him or her of being an incompetent English speaker. So this view will have a lot of problems to resolve until it can be worked out into a plausible view.

3.4.2. *The Logic of Proper Names*

Related to the question whether names have a lexical meaning is the question of which kind of *logical* inferences are licensed by proper names. Consider an utterance of the following sentence:

(1) Inma is a mathematician.

What can one logically infer from (1) – that is, what can one infer from (1) merely in virtue of what the words in (1) mean? For instance, here is a fairly uncontroversial case of something that logically follows from (1):

(2) Someone is a mathematician.

And here is a fairly uncontroversial case of something that does not logically follow from (1), even if (1) cannot be truly uttered without the following being also true:

(3) The speaker of (1) can speak some English.

But between the two, there are cases in which the answer is not as clear as for (2) or (3). For instance, is the following something that follows from (1) merely in virtue of what the words in (1) mean?

(4) Someone is called 'Inma'.

If your answer is 'yes', then you are on the side of the metalinguistic view, while if your answer is 'no', then you have a good reason to reject it.

I do not know the right answer to the question whether (4) logically follows from (1) – if there *is* a right answer, to begin with. What I know is that individual constants in FOL, with which we often "translate" proper names, do not support the inference from (1) to (4). For example, if you accept the metalinguistic thesis, and if you accept that logical inference is inference in virtue of what the words mean, then you should accept the following as a logical consequence of (1):

(5) There are two individuals, Inma and the name 'Inma', and the first is a
    bearer of the second.

But, of course, in FOL, a sentence of the form 'F(c)' (where F is a one-place predicate and c an individual constant), does not entail '$\exists x \exists y (x \neq y)$'.[82]

3.4.3. *Proper Names and* What Is Said

The third set of empirical data on which we should test our theories of proper names concerns their contribution to what is said. The mainstream view holds that what names contribute to what is said is their reference. But we saw that the same thesis with respect to indexicals proved to be problematic, so the question becomes whether, when it comes to proper names, we have reasons to rehabilitate the mainstream view, rather than seek an alternative.

---

[82] $\exists x \exists y (x \neq y)$ follows from (5) on the uncontroversially true assumption that a person's name is distinct from the person so-named. The argument presented in this subsection comes from Predelli (ms.). Predelli also sets out to explain why it is (normally) uninformative to tell someone that Inma is called 'Inma', and also, how you can explain on pragmatic grounds that any time that you assent to an utterance of (1), you (normally) know that (4) is true as well.

Consider the following Hesperus-Phosphorus kind of case. Suppose that Inma knows Prof. Cheng "under two guises": she knows her as the professor with whom she signed up for a class next term, as she knows her simply as Patricia, a friend of Tarek. Of course, she does not realize that Patricia is Prof. Cheng. Consider:

(6) Patricia will join us for dinner. (said by Tarek)

(7) Prof. Cheng will join us for dinner. (said by Maria)

On the referentialist view, Tarek in (6) and Maria in (7) have said the same thing. But take Inma's point of view. For her, (7) will still be informative. So, one might say, if Inma really understood what Tarek said in (6), and if she similarly understood what Maria said in (7), then she would have known that they said the same thing, hence (7) would be uninformative to her. The assumption that (6) and (7) say the same thing leads, then, to the implausible conclusion that Inma did not understand either what Tarek said or what Maria said. What's more, if Inma herself uses the names 'Patricia' or 'Prof. Cheng', we would need to conclude that she does not know what she is saying. And upon reflection, we would see that most times people would be saying things that they do not understand, which is an implausible view to hold.

Let us now turn to a different kind of case. Suppose that Tarek says to Inma:

(8) I've always wanted to visit Georgia.

A few days later, the following dialogue takes place:

(9) Maria: I've always wanted to visit Georgia.

(10) Inma: Tarek said that, too.

Suppose that it is common knowledge between Inma and Maria that Tarek has never heard of Maria, so that only the sloppy (or *de se*) reading is available (cf. 2.5.1). Then, of course, we are tempted to say that Inma's report is true. But suppose now that the Georgia that Tarek wants to visit is Georgia the ex-

republic of U.S.S.R., while the Georgia that Maria wants to visit is Georgia the state of U.S.A. What should we then say about Inma's report?

The first option is to say that Inma's report is false, but not because Tarek and Maria would have said different things. Rather, they said the same thing, but Inma's report is false because of an incorrectly resolved presupposition. Following tenet (f) of our theory (cf. 2.7), Inma, who is the reporter, will be presupposing that the reportee, that is, Tarek, was talking of the same thing as Maria, hence of the same Georgia, and said the same thing about it. But Tarek was talking of something different (of that other Georgia), hence he did not say the same thing as Maria about the Georgia that she was talking about. He did say the same thing as Maria, but about a different Georgia.

To see if this solution works, we need to see what happens when we "fix" the presupposition. Consider:

(11) Maria: I've always wanted to visit Georgia.

(12) Inma: Tarek said that, too.

(13) Maria: Yes, he said that, but he was talking of Georgia the ex-republic of U.S.S.R. Me, I'm talking of Georgia, U.S.A.

If we are happy with Maria's reply, then the solution seems to work fine.

The second option is, of course, to go referentialist. One would then say that Inma's report is false because Tarek and Maria actually said different things:

(14) Maria: I've always wanted to visit Georgia.

(15) Inma: Tarek said that, too.

(16) Maria: No, what he said is that he always wanted to visit Georgia the ex-republic of U.S.S.R. What I said is that I've always wanted to visit Georgia, U.S.A.

If we think that the correct reply for Maria is (16), rather than (13), then we should opt for the referentialist solution. At this stage, though, we need not decide which solution is better.

3.4.4. The Referentialist View

I take it to be unnecessary to lay down the referentialist view on proper names. Rather, what I want to do in this subsection is show that accepting my view on indexicals does not commit you to giving up the referentialist view on names: the two are mutually consistent. However, I will point to possible tensions between the motivations behind the two views.

The easiest way to see that the two views are consistent is by looking at the formal framework. The formal semantics of which I have given bits and pieces throughout Chapter 1 and which is laid down in Appendix 1 is quantified modal logic, in which (non-modal and non-temporal) indexicals are taken care of by means of free variables, plus the lexically encoded conditions indexed to the designated index. The point is that there may well be individual constants in this framework. Individual constants have their interpretation fixed by the interpretation function, determined by the interpretation structure. The interpretation of an individual constant is a mapping from world-time pairs to individuals. If it is a *constant* mapping (i.e., if it assigns the same individual to every world-time pairs), then the constant is a rigid designator. Suppose that you decide to translate a name such as 'Tarek Al'Mouhari' by such a rigid individual constant. That is entirely possible within my framework. To do so precisely amounts to endorsing the referentialist view for proper names. So, what this shows is that my view on indexicals and the referentialist view of proper names are perfectly compatible.

Now, if you ask me whether I would opt for the referentialist view, I probably would not. The reason is that some of the things that I said to lend plausibility to my view on indexicals do not fit well with the spirit of the referentialist view, even when the latter is confined to names. One of the ideas

(in 1.2) was that when I say "She is dangerous" to communicate to you that Suleïka is dangerous, I am referring directly to Suleïka and saying *of her* that she is dangerous, and the condition of being female, lexically encoded in 'she', merely helps you figure out that it is Suleïka that I am talking about. Now suppose that, in the same scenario, there are several women around, so the gender constraint is no longer very helpful. But if I know that you know that Suleïka's name is 'Suleïka', then I could just tell you "Suleïka is dangerous". It would be nice to explain what happens here along more or less the same lines as in the case of what happened when I used the indexical 'she'. We might want to say that here, too, I was directly referring to Suleïka, and that my use of her name was merely heuristic, to help you figure out that I was talking about her. Now, though this is not incompatible with the idea that the semantic content of the name 'Suleïka' is Suleïka herself, it is also true that if we endorsed the referentialist view on names, we would get something of an asymmetry between the two cases, an asymmetry that we might want to avoid.

Another reason why I would not go for the referentialist view on names is that I think that when it comes to the question of what names contribute to what is said, there is more evidence against than for referentialism. Not only because of the Hesperus-Phosphorus cases, but even in the case of Georgia, where I believe that the falsity of the report in (10) is due to an incorrectly resolved presupposition, rather than to some difference in what was actually said in (8) and (9). Also, I am perfectly happy with Maria's reply in (13).

3.4.5. *The Metalinguistic View*
The metalinguistic view holds that proper names have meanings, and that the meaning of a given name is obtained through the following schema:

(the metalinguistic thesis) The lexical meaning of 'N' is "the bearer of 'N'."

The weaker version of the metalinguistic view, held, for instance, by Recanati runs, sees the metalinguistic thesis as a thesis about the *character* of a name,

but when it comes to the *content*, the view holds, with Kripke and Kaplan, that the proper name contribute its reference, and nothing but its reference. Others, e.g. Geurts (1997), hold a stronger version of the metalinguistic view, since they take a proper name to be just a shorthand for the definite description "the bearer of [that name]." It is interesting that both versions are vulnerable to the argument from unwelcome logical consequence (cf. 3.4.2. and Predelli (ms.)).

If I were to hold the metalinguistic view, here is what it would look like. We would not need to find some unique definite description that would capture the lexical conditions associated with the proper name. Descriptions such as "bearer of the name 'David'" are merely approximations to the lexical meaning of the name 'David'. But the lexical meaning, hence the semantic content, of a name such as 'Tarek' will be exactly like the semantic content of any one-place predicate, namely, a function from world-time pairs to sets of individuals (sets that contain all and only Davids). As in the case of indexicals, the semantics of proper names would require double indexing. Consider:

(17) It is possible that Tarek should not bear the name 'Tarek'.

We want (17) to turn out true, which is why the condition lexically associated with the name is to be indexed to the designated index and ignore the current index of evaluation. I take it that all of this is fairly straightforward.

Here is how a sentence such as "Suleïka is dangerous" would get analyzed on this version of the metalinguistic view:

$S, w_1, t_1, w_2, t_2, f \vDash$ dangerous(Suleïka)   iff

$S, w_1, t_1, w_2, t_2, f \vDash$ dangerous(x) and $S, \boxed{w_2, t_2,} w_2, t_2, f \vDash$ is-a-Suleïka(x)

Of course, the speaker and the hearer would need to pragmatically convene on taking f to be the assignment that assigns Suleïka to x. As in the case of indexicals, the lexical meaning of the name is helpful for communication, because it constrains the range of assignments that will make the sentence true. The hearer will know that when I tell him "Suleïka is dangerous" I am

not talking about Inma, because she is not "a Suleïka" (unless 'Suleïka' is one among her many names).

Now, if you ask me whether I would opt for the metalinguistic view, I might say 'yes', but it would be a conditional 'yes': yes, if there is nothing better. The reason why I am not so enthusiastic about this view is twofold. On the one hand, proper names certainly do not have the same kind of lexical meaning as other words do, like 'pencil' or 'bench' or 'she' or 'today' or 'necessarily'. We might say, with Recanati, that they have some special kind of lexical meaning, but we might also plausibly prefer to say that they have no lexical meaning at all. Secondly, going back to 3.5.2, I am inclined to say that the inference from (1) to (4) is not a logical inference, but only a pragmatic inference. Now, that is a way of saying that the condition in the metalinguistic thesis is not to be handled at the level of semantics. Those are, then, the reasons why I would not go for the metalinguistic view either.

3.4.6. The "No Semantics" View

My main proposal has been that the semantic content is lexical meaning – for indexical and non-indexical expressions alike. But I have ended the previous subsection suggesting that proper names do not have any lexical meaning. It follows that proper names have no semantic content at all.

Though this might sound as a *reductio* of one of the two hypotheses, it is not intended to be. I actually find the idea that proper names have no semantic content attractive, and what I would like to do now is give some plausibility to that idea. The first question is, if proper names have no semantics, what are they? The suggestion would be that names are merely *pragmatic* devices that are used by the speaker to help the hearer figure out to whom she is referring. For instance, if I use the name 'Tarek', that will help you figure out that I am referring to Tarek – of course, only if you, too, are competent with this pragmatic device, that is, if you know who Tarek is. Sometimes, as with other such devices, the name will not be a sufficient clue. If I tell you that David is

coming, you might ask, *which David*? David Kaplan, David Hills, David Israel... The name is merely a helpful clue for figuring out who is being referring, but in many cases, further contextual clues will be required as well.

Now, one might find it worrisome that language might contain elements that are merely pragmatic devices and have no semantics. But proper names would not be alone in that respect. Consider exclamatives: the expression 'ouch' is part of the language, and so is 'wow', and the like. Yet we do not expect those expressions to have a semantics. Similarly, words like 'frankly' in "Frankly, she is a genius" are often seen as devoid of semantic content. Of course, with proper names the situation might be more subtle. For one thing, names are used in the formation of sentences, they occupy the place of the grammatical subject in the same way in which pronouns, definite descriptions, or quantifiers do. One could then worry that stripping proper names of any semantics would be incompatible with compositional semantics.

If we think of the semantic content along the lines of my proposal from Chapter 1, then we should not find the "no semantics" idea as problematic as we might have otherwise. Recall our paradigmatic case of direct reference: we are staring at Rembrandt's *Night Watch*, and I say:

(18) Impressive!

Now consider a case in which I similarly want to say of that very painting that it is impressive, but suppose that though we are in front of the painting, you are actually going through a catalogue with many other impressive paintings. In such a situation, I would say:

(19) *The Night Watch* is impressive!

The idea is, then, that in (19) I am referring directly to the painting – just as in the case of (18) – and I want to communicate something about that painting to you, but because Rembrandt's *Night Watch* is not the most salient object, and you do cannot guess in advance that I would be talking about that painting, I

need to help you realize that I am referring to *The Night Watch*. So I use the painting's name. This name does not contribute anything to the semantic content, it only helps you figure out which painting I am talking about. What the name does is not very different from what I would do if I had pointed at the painting. The pointing gesture is clearly a pragmatic device. So might be the name.

There remains the question of what the semantic system does when you give it as input a sentence that contains a name. Of course, there is the option of cleaning up the sentence of all pragmatic elements before you feed it into the semantic system. That is presumably what you do with exclamatives like 'ouch', as well as with adverbs such as 'frankly'. Or else, you can leave those in the input, but teach the system to ignore them.

However, if we ask ourselves what a possible truth clause for a sentence containing a name might look like on the "no semantics" view, here is what immediately springs to mind:

$$S, w_1, t_1, w_2, t_2, f \vDash \varphi(\text{Suleïka}) \quad \text{iff} \quad S, w_1, t_1, w_2, t_2, f' \vDash \varphi(x)$$
$$\text{where } f'(x) = \text{Suleïka and } f' \text{ is like } f \text{ otherwise}$$

While this is perfectly intelligible for anyone even remotely familiar with formal logic, I want to stress that the proposed truth clause is not innocent: it short-circuits the passage through the structure of interpretation. The truth clause roughly tells you, if you come across the name Suleïka, go and get Suleïka directly, and assign her to the variable x. One thing that this entails is that for every name in the object-language, there would be a corresponding name in the meta-language. Without this, it would be impossible to even write down the proposed truth clause. So, on this proposal, the interpretation of a name is very different from the interpretation of an individual constant:

$$S, w_1, t_1, w_2, t_2, f \vDash \varphi(c) \quad \text{iff} \quad S, w_1, t_1, w_2, t_2, f' \vDash \varphi(x)$$
$$\text{where } f'(x) = I^S(\text{'c'}), \text{ and } f' \text{ is like } f \text{ otherwise}$$
$$(I^S \text{ is the interpretation function of structure } S)$$

In conclusion, if you ask me if I would opt for the "no semantics" view, I might say 'yes', though this would again be a conditional 'yes'. Yes, provided that the view may be worked out in a satisfactory way. All that I have done here is to outline the view and to show that it is not as implausible as you might have thought at first. But working out the details of this view is far beyond the scope of this work. The central topic of this work were indexicals, their semantics, and what is said by speakers uttering sentences containing indexicals. And there, my proposal is that the lexical meaning = the semantic content = what is said.

——

end of Chapter 3

——

Appendix 1

The Formal Framework

In what follows, I lay down the formal framework. Let me stress, though, that this is just *a* possible semantic framework; there are other ones that might equally well, if not even better, account for the formal aspects of my proposal.

A.1.1. The Formal Language and Its Syntax

- The language of quantified modal tense logic with indexicals:

  p, q, $future^0$, $past^0$, $tomorrow^0$, $yesterday^0$, P^1x, Q^1she_x, $R^2c_1I_y$, $R^3he_xc_2he_y$,

  $\neg \varphi$, $\varphi \rightarrow \psi$, $\varphi \wedge \psi$, $\Diamond \varphi$, $\Box \psi$, actually ψ, sometime-future ψ, always-past φ,

  now ψ, the-day-after φ, $\forall x \varphi x$, $\exists y \psi y$...

Remarks. Among the propositional atoms, such as p and q, that are part of the non-logical vocabulary, we will have dates and days in the week; for example, '$Friday^0$', '$January\text{-}20\text{-}2007^0$', etc. On the other hand, '$future^0$', '$past^0$', etc., are propositional atoms, but are also part of the logical vocabulary. See 1.9.1. for motivations for this formal treatment and for discussion.

A.1.2. Structures of Interpretation and Models

- Kripkean structures for quantified modal tense logic:

  $S := (U, W, T, R^W, \leq, \geq, \vartriangleleft, I^S)$, where $R^W \subseteq W \times W$, $\leq, \geq, \vartriangleleft \subseteq T \times T$, and

  $I^S('c_1') := W \times T \mapsto U$, $I^S('p') := W \times T \mapsto \{T, F\}$, $I^S('P^1') := W \times T \mapsto \wp(U)$,

  $I^S('R^2') := W \times T \mapsto \wp(U \times U)$, etc.

Remarks. 'X↦Y' stands for a function that sends the elements of X to elements of Y. So, for instance, given any pair (w,t), I^S('P') will return some subset of U. R^W is the accessibility relation on W, and ≤ and ≥ are partial orders on T, used in interpreting future and past modalities. The relation ◁ is a novelty: it is a relation on T such that, intuitively, t◁t' if and only if t' is, or falls on, the day after t; we need this relation to provide the truth clauses for 'tomorrow' and 'the day after'.

- Models for double-indexed semantics:

  $M := (S, w_1, t_1, w_2, t_1, f)$, where $w_1, w_2 \in W^S$, $t_1, t_2 \in T^S$,

  and f is an assignment of values to the free variables.

Remarks. By defining the notion of a model as requiring two, rather than one, world-values and time-values, we may already foresee that the resulting logic will be different from Kaplan's. For instance, 'actuallyψ↔ψ' will not be valid, while the rule of necessitation will hold. See 1.5.4. for discussion.

A.1.3. Semantics

- Truth clauses for atomic sentences without indexicals

  $S, w_1, t_1, w_2, t_2, f \vDash p$ iff $\quad\quad\quad I^S('p')(w_1, t_1) = \top$

  $S, w_1, t_1, w_2, t_2, f \vDash R^2xc$ iff $\quad\quad (f(x), I^S('c')(w_1, t_1)) \in I^S('R')(w_1, t_1)$

- Truth clauses for non-indexical expressions

  $S, w_1, t_1, w_2, t_2, f \vDash \neg\, \varphi$ iff $\quad S, w_1, t_1, w_2, t_2, f\ \nvDash \varphi$, when φ contains no indexical pronouns (otherwise see below)

  $S, w_1, t_1, w_2, t_2, f \vDash \varphi \to \psi$ iff $\quad S, w_1, t_1, w_2, t_2, f \vDash \neg\, \varphi$

  $\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad$ or $S, w_1, t_1, w_2, t_2, f \vDash \psi$

  $S, w_1, t_1, w_2, t_2, f \vDash \varphi \wedge \psi$ iff $\quad S, w_1, t_1, w_2, t_2, f \vDash \varphi$

  $\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad$ and $S, w_1, t_1, w_2, t_2, f \vDash \psi$

  $S, w_1, t_1, w_2, t_2, f \vDash \Diamond \varphi$ iff for some w' s.t. $w_1 R w'$: $S, w', t_1, w_2, t_2, f \vDash \varphi$

  $S, w_1, t_1, w_2, t_2, f \vDash \Box \varphi$ iff for every w' s.t. $w_1 R w'$: $S, w', t_1, w_2, t_2, f \vDash \varphi$

  $S, w_1, t_1, w_2, t_2, f \vDash$ sometime-future φ iff $\quad$ for some t' s.t. $t_1 \leq t'$:

  $\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad S, w_1, t', w_2, t_2, f \vDash \varphi$

$S, w_1, t_1, w_2, t_2, f \vDash$ always-past φ iff   for every t' s.t. $t_1 \geq t'$:

                                                       $S, w_1, t', w_2, t_2, f \vDash \varphi$

$S, w_1, t_1, w_2, t_2, f \vDash$ the-day-after φ iff   for some t' s.t. $t_1 \lhd t'$:

                                                       $S, w_1, t', w_2, t_2, f \vDash \varphi$

$S, w_1, t_1, w_2, t_2, f \vDash \forall x \varphi x$ iff   for every assignment f' that is like f except

                                          at most for x: $S, w_1, t_1, w_2, t_2, f' \vDash \varphi x$

$S, w_1, t_1, w_2, t_2, f \vDash \exists x \varphi x$ iff   for some assignment f' that is like f except

                                          at most for x: $S, w_1, t_1, w_2, t_2, f' \vDash \varphi x$

- Truth clauses for temporal and modal indexicals

$S, w_1, t_1, w_2, t_2, f \vDash$ actually φ iff $S, \mathbf{w_2,} t_1, w_2, t_2, f \vDash \varphi$

$S, w_1, t_1, w_2, t_2, f \vDash$ now φ iff $S, w_1, \mathbf{t_2,} w_2, t_2, f \vDash \varphi$

$S, w_1, t_1, w_2, t_2, f \vDash$ tomorrow   iff $t_2 \lhd t_1$

$S, w_1, t_1, w_2, t_2, f \vDash$ yesterday   iff $t_1 \lhd t_2$

$S, w_1, t_1, w_2, t_2, f \vDash$ future   iff $t_1 \geq t_2$ and $t_1 \neq t_2$

$S, w_1, t_1, w_2, t_2, f \vDash$ past   iff $t_1 \leq t_2$ and $t_1 \neq t_2$

- Truth clauses for indexical pronouns

$S, w_1, t_1, w_2, t_2, f \vDash \varphi(I_x)$ iff $S, w_1, t_1, w_2, t_2, f \vDash \varphi(x)$ and

                                           $S, w_2, t_2, w_2, t_2, f \vDash$ speaker(x)

$S, w_1, t_1, w_2, t_2, f \vDash \neg \varphi(I_x)$ iff $S, w_1, t_1, w_2, t_2, f \vDash \neg \varphi(x)$ and

                                           $S, w_2, t_2, w_2, t_2, f \vDash$ speaker(x)

$S, w_1, t_1, w_2, t_2, f \vDash \psi(you_y)$ iff $S, w_1, t_1, w_2, t_2, f \vDash \psi(y)$ and

                                           $S, w_2, t_2, w_2, t_2, f \vDash$ addressee(y)

$S, w_1, t_1, w_2, t_2, f \vDash \neg \psi(you_y)$ iff $S, w_1, t_1, w_2, t_2, f \vDash \neg \psi(y)$ and

                                           $S, w_2, t_2, w_2, t_2, f \vDash$ addressee(y)

$S, w_1, t_1, w_2, t_2, f \vDash \varphi(she_x)$ iff $S, w_1, t_1, w_2, t_2, f \vDash \varphi(x)$ and

                                           $S, w_2, t_2, w_2, t_2, f \vDash$ female(x)

$S, w_1, t_1, w_2, t_2, f \vDash \neg \; \phi \,(\text{she}_x)$ iff $\quad$ $S, w_1, t_1, w_2, t_2, f \vDash \neg \; \phi \,(y)$ and

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ $S, w_2, t_2, w_2, t_2, f \vDash \text{female}(y)$

$S, w_1, t_1, w_2, t_2, f \vDash \psi(\text{he}_y)$ $\quad$ iff $\quad$ $S, w_1, t_1, w_2, t_2, f \vDash \psi(y)$ and

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ $S, w_2, t_2, w_2, t_2, f \vDash \text{male}(y)$

$S, w_1, t_1, w_2, t_2, f \vDash \neg\, \psi(\text{he}_y)$ $\quad$ iff $\quad$ $S, w_1, t_1, w_2, t_2, f \vDash \neg\, \psi(y)$ and

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ $S, w_2, t_2, w_2, t_2, f \vDash \text{male}(y)$

$S, w_1, t_1, w_2, t_2, f \vDash \varphi(\text{this}_z)$ $\quad$ iff $\quad$ $S, w_1, t_1, w_2, t_2, f \vDash \varphi(z)$ and

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ $S, w_2, t_2, w_2, t_2, f \vDash \text{salient}(z)$

$S, w_1, t_1, w_2, t_2, f \vDash \neg\, \varphi(\text{this}_z)$ $\quad$ iff $\quad$ $S, w_1, t_1, w_2, t_2, f \vDash \neg\, \varphi(z)$ and

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ $S, w_2, t_2, w_2, t_2, f \vDash \text{salient}(z)$

Remarks. The truth clauses for the indexical pronouns crucially involve non-logical vocabulary: to determine whether $S, w_2, t_2, w_2, t_2, f \vDash \text{speaker}(x)$, one needs to see whether $f(x)$ belongs to $I^S(\text{'speaker'})(w_2, t_2)$, which only gets determined by the structure of interpretation. Also, note that for each pronoun, we have given a *pair* of truth clauses, depending on whether the indexical lies in the scope of the negation operator or not. Though the result might lack elegance, it is empirically adequate; see 1.7.2. (fn. 32) for discussion.

A.1.4. Logic

- Sentence ψ is logically valid iff for any structure S, any $w_1, w_2 \in W^S$, any $t_1, t_2 \in T^S$, and any assignment function f: $S, w_1, t_1, w_2, t_2, f \vDash \psi$

- ψ is a logical consequence of φ iff for any S, w_1, w_2, t_1, t_2 and f such that $S, w_1, t_1, w_2, t_2, f \vDash \varphi$, we also have: $S, w_1, t_1, w_2, t_2, f \vDash \psi$

____

end of Appendix 1

____

Semantic Equivalence Results

In what follows, I will show that from the viewpoint of semantics, it does not make that much of a difference whether we use a variable for some parameter that truth depends upon, or rather, add another "primitive" parameter in the parametrization of truth. For the sake of simplicity, I will show the semantic equivalence for the case of the taste parameter (cf. 3.2). For the sake of clarity, I will ignore the time parameter, and since we will not be interested in 'actually' or any other indexical, I will just use single-indexed semantics.

A.2.1. The Result Towards Which We Are Heading

I am going to show that the classical, implicit argument semantics (CS), in which the taste parameter is seen as an argument to the taste predicate, and the relativist semantics (RS), in which it is seen as a parameter of the index of evaluation, on a par with the possible world, are much more similar than is commonly believed. I will show that the two semantic theories are equivalent by defining a bi-directional translation procedure T for which the following holds.

Let S_c and S_r be respectively sentences in the languages of CS and of RS, let f_1, f_2 be assignments of values to the free variables, w a world of evaluation, and u a "taste" of evaluation. Then:

- S_r is true with respect to f_1, w and u <u>iff</u> $T(S_r)$ is true with respect to f_1^T and w, where assignment f_1^T is defined in terms of f_1 and u.

- S_c is true with respect to f_2 and w <u>iff</u> $T(S_c)$ is true with respect to f_2, w and u^T, where u^T is a taste value obtained directly from f_2;

The method that I am using is classic, and the result obtained, *qua* a formal result, should be old news to anyone familiar with modal logic.[83] But I still take it to be worthwhile to lay down the result, as it appears to be ignored by almost all the protagonists the debates on taste predicates, epistemic modals, knowledge reports, and the like.

A.2.2. *The Classical Implicit Argument Semantics for Taste Predicates*

The target for CS is a small fragment of English, containing words like 'delicious' or 'lousy', demonstrative and 3rd person pronouns, names, basic quantifiers (everyone, someone) and basic modalities (necessarily, possibly). The formal language is just the language of quantified modal logic, but the semantics is as simple as it can get: we ignore accessibility relations and keep the universe fixed across worlds. Also, even though we have demonstratives in our target language, we will translate them as free variables, so that we do not even need two-dimensional semantics. The only novelty is that we have one distinguished variable, x_T, which we use for the implicit argument associated with any taste predicate. For the sake of simplicity, let the values assignable to this variable be just ordinary individuals. The main task of what we informally think of as context is to provide an assignment of values to free the variables, and to supply the world at which the sentence is evaluated for its truth value.

Language. $L_{CS}:= P_1^0, ..., P^1_1,..., P^n_1,...; x_T, x_1, x_2, ...; c_1, c_2, ...; \neg , \wedge, \diamondsuit, \square, \forall, \exists$.

---

[83] It essentially derives from the equivalence between modal logic S5 and monadic predicate logic.

Syntax. The standard recursive definition of well-formed formulae. The only somewhat unusual constraint is that for any atomic sentence, variable x_T may occur at most once, and only in the very last position.[84]

Semantics. S is a structure iff$_{df}$ it is of the form (U, W, I), where U is the universe, W is the set of possible worlds, and for every P^n_i, I('P^n_i') will be a function that maps possible worlds to sets of n-uples of individuals.

The truth value of a sentence is defined at a model. A model is a triple of the form (S, w, f), where S is a structure, w a designated world, and f an assignment of values to the free variables. The truth definition is recursive and entirely standard, so I will not bother to spell it out.

To see how this works, take the sentence "This is delicious", as used to express one's own taste. Its translation will be DELICIOUS x_1 x_T , so, for instance, if Tarek utters this sentence while referring to the chocolate cake, the assignment of values relevant to the truth of his utterance will send x_1 to the chocolate cake and x_T to Tarek, and if Inma utters the same sentence in reference to a certain cookie, we will normally assign that cookie to x_1 and Inma to x_T. And if Tarek, or anyone else for that matter, uses the sentence to make a generic claim, meaning that a given chocolate cake is delicious *for everyone*, then the translation is going to be $\forall x_T$ DELICIOUS x_1 x_T, and the relevant assignment will send x_1 to that cake.

A.2.3. The Relativist Semantics for Taste Predicates

What the classical implicit argument view translates by a 2-place predicate one of whose argument is occupied by the taste variable x_T, the relativist will translate by a 1-place predicate. On the semantic side, though, the interpretation of such a one-place predicate is not just a

---

[84] For example, $P^2 x_T x_T$ and $P^2 x_T x_1$ are not well-formed, while $P^2 x_1 x_1$ or $P^2 x_1 x_T$ are. We need the "at most once" constraint for the equivalence results. However, the language thus constrained seems fine for modeling natural language, because, at least in English, we do not have atomic expressions that involve *two* taste parameters. As for the constraint that x_T always occurs at the end, it is only there to make it easier to define the translation between the two formal languages, but nothing important hinges on it.

mapping from possible worlds to sets of individuals. Rather, it is a mapping from *pairs* (possible world, individual) to sets of individuals. The framework that follows is as in Lasersohn (2005), except for some minor differences.[85]

Language. L_{CS}:= P_1^0, ..., P_1^1,..., P_1^n,...; x_1, x_2, ...; c_1, c_2, ...; $\neg$, $\wedge$, $\Diamond$, $\Box$, $\forall$, $\exists$, ∎.

Syntax. The standard rules for well-formed formulae, which I will not bother to spell out, plus the rule: if φ is a formula, then so is ∎φ.

Semantics. S is a structure iff$_{df}$ it is of the form (U, W, I), where U is the universe, W is the set of possible worlds, and I is the interpretation function such that I('P_i^n') maps pairs (possible world, individual) to sets of n-uples of individuals.

The truth of a sentence is defined at a model, which is now a quadruple (S, w, u, f), where S is a structure, w a designated world, u a designated individual (who serves as the value for the taste parameter), and f an assignment of values to free variables. The definition of truth is recursive and standard. The only novelty is the truth clause for ∎:

S, w, u, f $\vDash$ ∎φ iff   for every u'$\in$U: S, w, u', f $\vDash$ φ

The syntax and the semantics for **L$_{RS}$** are perfectly standard. As in the case of CS, we take the usual semantics for modal predicate logic. It is a

---

[85]Among those minor differences, a major one is that we, unlike Lasersohn, have a universal operator on the taste parameter, ∎. Another is that Lasersohn also has in the formal language a class of predicate modifiers 'for c' (where c is a constant, e.g. one that translates a proper name like 'Tarek'), which are meant to translate complex expressions such as 'tasty for Tarek'. On the semantic side, 'for Tarek' works as a rigidifier: it makes the semantic value of 'tasty for Tarek' a constant function in the taste parameter, whose value, for any other individual, is the same as the value that 'tasty' alone takes at Tarek. In the implicit argument approach, the expression 'for c' simply makes the second argument of 'tasty' explicit. This argument, when *implicit*, is taken account of through the taste variable x_T. E.g. while 'tasty' is translated by Px_1x_T, 'tasty for Tarek' is translated as Px_1c_1, where the interpretation of c_1 is Tarek. I will leave the 'for'-construction out of consideration. It is a boring but easy exercise to construe extensions of CS and RS that make room for 'for' while remaining equivalent. Finally, Lasersohn's framework is cast within a Kaplanian framework, hence there is also a time parameter and a context parameter, both of which we are ignoring.

two-dimensional modal logic, since in addition to the possible world parameter w, there is the taste parameter u, with respect to which truth is recursively defined. The only expression that requires recursion on the taste parameter is the *universal taste operator* ■, which we are going to use in translating *generic* taste claims.[86] So for instance, is "This is delicious" is used by a person to express people's taste in general, in the sense of "delicious for everyone", its translation will be ■DELICIOUS x_1, and the relevant assignment will send x_1 to the object referred to with 'this'. But if a person uses the same sentence to express her own taste, the translation will be DELICIOUS x_1. If Tarek uses the sentence, then, given an assigment of value to x_1, it will be evaluated for truth at the index whose taste parameter is set to value Tarek, and if it is Inma who utters, then the value will be set to Inma.

A.2.4. Equivalence

To show the equivalence between $\mathbf{L_{CS}}$ and $\mathbf{L_{RS}}$, we need to define a suitable translation between the two languages. This is made easy by the fact that $\mathbf{L_{CS}}$ has a distinguished variable x_T (and also, that we do not have any accessibility relations on the taste dimension).[87] Here is the proposed translation T that takes $\mathbf{L_{RS}}$-formulas to $\mathbf{L_{CS}}$-formulas:

$T(\mathrm{P}c_1,...,c_m, x_1,..., x_n) = \mathrm{P}^T c_1,..., c_m, x_1,..., x_n, x_T$

$T(\neg\ \varphi) = \neg\ T(\varphi)$

$T(\varphi \wedge \psi) = T(\varphi) \wedge T(\psi)$

$T(\forall x_i \varphi x_i) = \forall x_i T(\varphi x_i)$

---

[86]It is not clear whether Lasersohn would acknowledge at all the generic readings of claims about taste – readings on which claiming that something is fun is to claim that it is fun for people in general. If he would, my hunch is that, rather than having a universal taste operator ■, he would use the construction 'for x' and bind the variable x with a regular ∀. Surprisingly, Lasersohn (2005) gives a language without quantifiers, and also, without any expression whose semantics requires recursion on the taste parameter.

[87] Even if we had accessibility relations on the taste parameter, that would only be a minor complication. In what follows, I rely on some known results from modal logic. See e.g. van Benthem (1983: 40).

$$T(\square\varphi) = \square T(\varphi)$$
$$T(\blacksquare\varphi) = \forall x_T T(\varphi).$$

Only the first and the last line are interesting. Remember that $\mathbf{L_{SR}}$ handles the dependence of taste predicates on the taste parameter by means of an additional parameter in the definition of truth. $\mathbf{L_{CR}}$, on the other hand, deploys no such parameter, but it has an additional argument place in every taste predicate. That is what the translation function T reflects. In fact, T "opens" a new argument place in every predicate, be it a taste predicate or not. For the latter, though, this argument will be idle.[88] What the first clause says is, roughly, that if P is a (n+m-place) predicate in the language $\mathbf{L_{RS}}$, then take a (n+m+1-place) predicate P^T in the language $\mathbf{L_{CS}}$, and use variable x_T in its last argument place. We also need to ensure that the structures of interpretation S of RS and S^T of CS will interpret two the predicates alike:

- if $(w, u, (u_1,..., u_{n+m})) \in I^S(P)$, then $(w, (u_1,..., u_{n+m}, u)) \in I^{ST}(P^T)$

What translation T does, too, in the last clause, is translate $\blacksquare$, the universal modal operator on the taste dimension, by the universal quantifier on the taste variable x_T.

In the other direction, our translation function T maps sentences of $\mathbf{L_{CS}}$ to sentences of $\mathbf{L_{RS}}$ as follows:

$$T(Pc_1...,c_m, x_1...,x_n,...,x_T) = P^t c_1...,c_m, x_1,...,x_n$$
$$T(Pc_1,...,c_m, x_1,...,x_n) = Pc_1,..., c_m, x_1,...,x_n$$
$$T(\neg\ \varphi) = \neg\ T(\varphi)$$
$$T(\varphi\wedge\psi) = T(\varphi)\wedge T(\psi)$$

---

[88] In this respect, there will be a difference between the "direct" translation of a taste-insensitive predicate in the language of CS, which is represented by a one-place predicate letter P, and the translation that we would get if we first translated it in the language of RS and then translated it into $\mathbf{L_{CS}}$ using the translation procedure T, since, then, it will be represented by a *two*-place predicate letter P^T. On the semantic side, though, the difference becomes unnoticeable.

$$T(\forall x_i \varphi x_i) = \forall x_i T(\varphi x_i)$$

$$T(\forall x_T \varphi x_T) = \blacksquare T(\varphi x_T)$$

$$T(\square \varphi) = \square T(\varphi).$$

Again, we must ensure that the structures of interpretation S of CS and S^T of RS will interpret the atomic predicates alike:

- if $(w, (u_1,..., u_{n+m}, u)) \in I^S(P)$, then $(w, u, (u_1,..., u_{n+m})) \in I^{S^T}(P^T)$;

- if $(w, (u_1,..., u_{n+m})) \in I^S(P)$, then $(w, v, (u_1,..., u_{n+m})) \in I^S(P)$, for all $v \in U$.[89]

The following results follow straightforwardly from the definitions of T, and may be verified by induction on the complexity of φ:

claim 1: $S, w, u, f \models_{RS} \varphi$ i f f $S^T, w, f^T \models_{CS} T(\varphi)$,

where $f^T(x_T) = u$ and f^T is otherwise like f

claim 2: $S, w, f \models_{CS} \varphi$ i f f $S^T, w, f(x_T), f \models_{RS} T(\varphi)$

____

end of Appendix 2

____

---

[89] Recall that in the language of CS, the taste argument is only associated with taste predicates. That is why we need the second clause in the translation procedure, as well as this constraint on interpretation structures: its task is to prevent the possibility of a taste-insensitive predicate from L_{CS} getting interpreted, in RS, by a function that varies in the taste parameter.

References

Bach, K. (1994). "Conversational Impliciture." *Mind and Language* 9: 124-62.

—— (2002). "Seemingly Semantic Intuitions." *Meaning and Truth*. Eds. J. Campbell et al. New York: Seven Bridges Press. 21–33.

Barwise, J. and R. Cooper, (1981). "Generalized quantifiers and natural language." *Linguistics and Philosophy* 4: 159–219.

—— and J. Perry (1999). *Situations and Attitudes*. Stanford: CSLI Publications.

van Benthem, J. (1983). *Modal Logic and Classical Logic.* Bibliopolis.

—— (1997). *Exploring Logical Dynamics.* Stanford: CSLI Publications.

Bianchi, C. (ed., 2005). *The Semantics/Pragmatics Distinction.* Stanford: CSLI Publications.

Blackburn, P. and J. Boss. (2003). "Computational Semantics". *Theoria* 18: 27-45.

Borg, E. (2004). *Minimal Semantics*. Oxford: Oxford University Press.

Büring, D. (2004). *Binding Theory*. Cambridge, MA: Cambridge UP.

Cappelen, H. and E. Lepore. (1997). "On an Alleged Connection Between Indirect Speech and the Theory of Meaning." *Mind and Language* 12: 278–96.

—— (2005). *Insensitive Semantics*. Oxford: Blackwell.

Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

—— (forthcoming). "Linguistic Communication and the Semantics/ Pragmatics Distinction." *The Semantics/Pragmatics Distinction.* Synthese Special Issue. Ed. I. Stojanovic.

Crossley, J. and L. Humberstone (1977). "The Logic of 'Actually'." *Reports on Mathematical Logic* 8: 11-28.

Dever, J. (1998). *Variables*. Ph.D. dissertation, University of California, Berkeley.

Donnellan, K. (1966). "Reference and Definite Descriptions." *Philosophical Review* 75: 281-304.

Evans, G. (1979). "Does Tense Logic Rest Upon a Mistake?" *Collected Papers*. Oxford: Clarendon Press (1985): 343–63.

—— (1980). "Pronouns." *Linguistic Inquiry* 11: 337 – 362.

Everett, A. (2000). *Saying, Samesaying, and What Is Said.* Ph.D. dissertation, Stanford University.

Feldman, R. (1980). "Saying Different Things." *Philosophical Studies* 38: 79–84.

von Fintel, K. and A. Gillies (ms). "CIA Leaks".

Frege, G. (1897). "Sense and Denotation." *The Frege Reader*. Ed. K. Beaney. Oxford: Blackwell (1997): 151-171.

—— (1899). "Logic." *The Frege Reader*. Ed. K. Beaney. 227-250.

—— (1917). "Thought." *The Frege Reader*. Ed. K. Beaney. 325-345.

Geach, P. (1962). *Reference and Generality*. Ithaca: Cornell University Press.

Geurts, B. (1997). "Good News for the Description Theory of Names." *Journal of Semantics* 14: 319-348.

—— (1999). *Pronouns and Presuppositions.* Amsterdam: Elsevier.

Graff, D. (2001). "Descriptions as Predicates." *Philosophical Studies* 102: 1-42.

Grice, P. (1989). *Studies in the Ways of Words*. Cambridge, MA: Harvard University Press.

Groenendijk, J. and M. Stokhof (1991). "Dynamic Predicate Logic." *Linguistics and Philosophy* 14: 39-1000.

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases.* Ph.D. Dissertation. University of Massachusetts.

—— (1990). "E-type Pronouns and Donkey Anaphora". *Linguistics and Philosophy* 13: 137-78.

—— (ms). "The First Person"

Hintikka, J. (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell: Cornell University Press

Kamp, H. (1971). "Formal Properties of 'Now'." *Theoria* 2: 227-73.

—— (1981). "A Theory of Truth and Semantic Representation" in *Formal Methods in the Study of Language*. Eds. J. Groenendijk et al. Amsterdam: Matematische Centrum. 277-321.

—— and U. Reyle (1993). *From Discourse to Logic*. Dordrecht: Kluwer.

Kaplan, D. (1977). "Demonstratives. An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals". Eds. J. Almog, J. Perry, and H. Wettstein. *Themes from Kaplan*. Oxford: Oxford UP (1989): 481–563.

—— (1989). "Afterthoughts". *Themes from Kaplan*. Eds. J. Almog et al. Oxford: Oxford UP. 564-614.

Keenan, E. and Y. Stavi. (1986). "A Semantic Characterization of Natural Language Determiners." *Linguistics and Philosophy* 9: 253-326.

King, J. (2003). "Tense, Modality, and Semantic Values." *Philosophical Perspectives* 17.

Kölbel, M. (2002). *Truth Without Objectivity*, Routledge.

Korta, K. and J. Perry (2006). "Three Demonstrations and a Funeral." *Mind and Language* 21: 166-186.

Kripke, S. (1979). "Speaker's Reference and Semantic Reference" *Contemporary Perspectives in the Philosophy of Language*. Eds. P. French et al. Minneapolis: University of Minnesota Press. 6-27.

—— (1980). *Naming and Necessity.* Cambridge: Harvard University Press

Lappin, S. (1996). "The Interpretation of Ellipsis." *The Handbook of Contemporary Semantic Theory*. Ed. S. Lappin. Oxford: Blackwell. 145–75.

Lasersohn, P. (2005), "Context Dependence, Disagreement, and Predicates of Personal Taste", *Linguistics and Philosophy* 28: 643-686.

Lasnik, H. (1989). *Essays on Anaphora*. Dordrecht: Kluwer.

Levinson, S. (2000). *Presumptive Meanings: The Theory of Generalized Conversational Implicatures*. Cambridge, MA: MIT Press.

Lewis, D. (1972). "General Semantics." *Natural Language Semantics*. Eds. D. Davidson and G. Harman. Dordrecht: Reidel. 169-218.

—— (1980). "Index, Context and Content." *Philosophy and Grammar*. Eds. S. Kanger and S. Ohman. Dordrecht: Reidel. 78–100.

—— (1981) "Attitudes *De Dicto* and *De Se*." *The Philosophical Review* 88: 513-43.

MacFarlane, J. (2003), "Future Contingents and Relative Truth." *The Philosophical Quarterly* 53: 321-36.

—— (forthcoming). "Relativism and Disagreement." *Philosophical Studies*.

—— (ms.) "Epistemic Modals Are Assessment Sensitive."

Neale, S. (1990). *Descriptions*. Cambridge: MIT Press.

Partee, B. (1973). "Some Structural Analogies between Tenses and Pronouns in English." *The Journal of Philosophy* 70: 601–9.

—— (1989). "Binding Implicit Variables in Quantified Contexts." *Chicago Linguistic Society* 25: 342-65.

Perry, J. (1993). *The Problem of the Essential Indexical and Other Essays*. Oxford UP

—— (1997). "Reflexivity, Indexicality and Names." *Direct Reference, Indexicality and Proposition Attitudes.* Eds. W. Kunne et al. Stanford: CSLI-Cambridge University Press.

—— (2001). *Reference and Reflexivity*. Stanford: CSLI Publications.

Predelli, S. (2002). "Intentions, Indexicals, and Communication". *Analysis* 62: 310-16.

—— (2005). *Contexts. Meaning,Truth, and the Use of Language*. Oxford: Oxford UP.

—— (2006). "The Problem with Token-Reflexivity". *Synthese* 148: 5-29.

—— and Stojanovic, I. (forthcoming). "Semantic Relativism and the Logic of Indexicals". *Relative Truth.* Eds. M. Kölbel and M. Garcìa-Carpintero. Oxford: Oxford UP.

—— (ms). "Socrates and 'Socrates' "

Prior, A. (1957). *Time and Modality*. Oxford: Oxford UP, 1957.

—— (1977). *Worlds, Times, and Selves*. Ed. K. Fine. London: Duckworth.

Recanati, F. (1993). *Direct Reference. Oxford:* Blackwell.

—— (2001). "What is Said." *Synthese* 128: 75–91.

—— (2004). *Literal Meaning*. Cambridge: Cambridge University Press.

Richard, M. *Time, Tense, and Reference*. Oxford: Oxford UP.

Salmon, N. (1986). *Frege's Puzzle.* Cambridge, MA: MIT Press.

Saul, J. (2002). "Speaker Meaning,What is Said, and What is Implicated." *Nous* 36: 228–48.

Schlenker, P. (2003). "A Plea for Monsters." *Linguistics and Philosophy 26*: 29-120.

Scott, D. (1970). "Advice in Modal Logic." *Philosophical Problems in Logic.* Ed. K. Lambert. Dordrecht: Reidel. 143-73.

Soames, S. (2001). *Beyond Rigidity: The Unfinished Agenda of "Naming and Necessity."* Oxford: Oxford UP.

Stainton, R. (2000). "The Meaning of "Sentences"." *Nous* 34: 441–54.

Stalnaker, R. (1978). "Assertion." *Syntax and Semantics 9*: *Pragmatics.* Ed. P. Cole. 315-332.

—— (1998). "On the Representation of Context." *JoLLI* 7: 3-19.

—— (1999). *Context and Content*. Oxford: Oxford UP.

Stanley, J. (2000). "Context and Logical Form." *Linguistics and Philosophy* 23: 391-434.

—— and Z. Szabo, (2000). "On Quantifier Domain Restriction*." Mind and Language* 15: 219-61.

—— and T. Williamson (1995). "Quantifiers and Context-Dependence." *Analysis* 55: 291-5.

Strawson, P. (1950). "On Referring." *Mind* 59: 320-334.

Stojanovic, I. (2003). "What to Say on What is Said." *Modeling and Using Context*. Eds. P. Blackburn et al. Berlin: Springer-Verlag, LNAI 2680. 300-313.

——. (2006). "What Is Said, Linguistic Meaning, and Directly Referential Expressions." *Philosophy Compass* 1: 10.111/j. 1747-9991.2006.00029. http://www.blackwell-compass.com/subject/philosophy/

—— (ed., forthcoming). *The Semantics/Pragmatics Distinction.* Special Issue of *Synthese.*

Strawson, P. (1950). "On Referring." *Mind* 59: 320-44.

Szabo, Z. (ed., 2006). *Semantics vs. Pragmatics*. Oxford: Oxford UP.

Taylor, K. (2007). "Misplaced Modification and the Illusion of Opacity." *Situating Semantics: Essays on the Philosophy of John Perry.* Eds. M. O'Rourke and C. Washington. MIT Press.

Turner, K. (ed., 1999). *The Semantics-Pragmatics Interface from Different Points of View*. Oxford: Elsevier.

Vallée, R. (ms). *Théories de la désignation.*

Ziff, P. (1972.) "What is Said." *Natural Language Semantics*. Eds. D. Davidson and G. Harman. Dordrecht: Reidel. 709-21.