PERSPECTIVAL THOUGHT. A PLEA FOR (MODERATE) RELATIVISM

François Recanati

Perspectival Thought

A plea for (Moderate) Relativism

François RECANATI

Institut Jean-Nicod

September 2006

Perspectival Thought
A plea for (Moderate) Relativism

Summary

BOOK 2 : EXPERIENCE AND SUBJECTIVITY

BOOK 3 : EGOCENTRICITY AND BEYOND

MODERATE RELATIVISM

Part One

The framework

Chapter 1

The distribution of content

Relativization

In modal logic, propositions are evaluated relative to possible worlds. A proposition may be true relative to a world *w*, and false relative to another world *w'*. A proposition whose truth-value varies across worlds is said to be *contingent* (as opposed to *necessary*). Similarly, in tense logic, propositions are evaluated relative to times. A proposition (e.g. the proposition that Socrates is sitting) may be true relative to a time *t*, and false relative to another time *t'*. A proposition that has this property is said to be *temporal* (as opposed to *eternal*).

Further applications of the relativization idea easily come to mind. The proposition that it is raining (at a given time, in a given world) is true relative to some places, and false relative to others. The proposition that one is a philosopher is true relative to some persons, and false relative to others. The proposition that spinach is delicious is true relative to some standards of taste, and false relative to others. The proposition that the treasure might be under the palm tree is true relative to some epistemic situations, and false relative to others. The proposition that John is tall is true relative to some standards of height, and false relative to others.

As I understand it, the relativization idea has two component sub-ideas, which I will call 'Duality' and 'Distribution'. Distribution presupposes Duality but, as we shall see, it is possible to accept Duality while rejecting Distribution.

1. [*Duality*] To get a truth-value, we need a circumstance of evaluation as well as a content to evaluate. (As Austin puts it, 'It takes two to make a truth'.)[1]

2. [*Distribution*] The determinants of truth-value distribute over the two basic components truth-evaluation involves : content and circumstance. That is, a determinant of truth-value, e.g. a time, is *either* given as an ingredient of content *or* as an aspect of the circumstance of evaluation.

Distribution

The distribution idea is apparent in the literature that stems from John Perry's work on unarticulated constituents.[2] According to Perry (1986b), if something is given as part of the situation which an utterance (or, for that matter, a mental representation) concerns, and against which it is evaluated, it does not have to be articulated in that representation. Thus Perry draws a distinction between 'It is raining here', which explicitly mentions a place, and 'It is raining', which leaves the place out of the picture. Perry describes the content of the latter not as a complete proposition but as a *propositional function*, true of some places and false of others. The place which actually determines the truth-value of the utterance is fixed not by the content of the utterance but by the situation which that utterance concerns (the situation the speaker manifestly intends to characterize). When I say 'It is raining here' the situation which my utterance concerns is typically more complex since it involves several places at once, between which a contrast is drawn (Recanati 1997, 2000b). In this case the place must be articulated and cannot be left out of the picture, since it is not independently fixed by the situation talked about.

The distribution idea also comes up in Kaplan's well-known argument for temporal propositions (Kaplan 1989 : 502-504). Kaplan's argument is based on the existence of temporal operators. The contents temporal operators operate on must be temporally neutral, Kaplan argued, for if they are not — if they are temporally-

---

[1] Austin 1971 : 124 n.

[2] See Part ÌX and especially Chapter 31 below.

specific — the temporal operators will be vacuous. A temporal operator specifies the time(s) with respect to which the proposition it operates on is to be evaluated. If the proposition itself specified a time, embedding the proposition under the temporal operator would have no effect whatsoever. Being already specified by the content to be evaluated, the time of evaluation would be fixed once for all and could no longer be shifted. So temporal operators must operate on temporally-neutral propositions — propositions which are true with respect to a time, and false with respect to another time, but which do not specify the time relative to which they are supposed to be evaluated.

To sum up, if a determinant of truth-value is an ingredient of content, it is pointless to attempt to provide it by manipulating the circumstance of evaluation (Kaplan). Being an ingredient of content, it is fixed 'before the encounter with the circumstance of evaluation'. (I borrow this metaphor from Kaplan, who uses it in a slightly different context.) Conversely, if it is given as an aspect of the circumstance of evaluation, it does not have to be articulated in the representation whose content is evaluated (Perry). The general principle which emerges is a principle of economy or optimality according to which a determinant of truth-value is either given as an ingredient of content or as an aspect of the circumstance of evaluation. The richer the circumstance, the poorer the content evaluated with respect to that circumstance can be ; and the richer the content, the poorer the circumstance. In particular :

• If the circumstance consists of a possible world only, the content must be a complete proposition (something that determines a function from possible worlds to truth-values).
• If the circumstance is richer and involves a time and a place in addition to a world, then the content can be less than fully propositional : it can be place- and time-neutral and determine only a propositional function (a function from place-time pairs to functions from possible worlds to truth-values, or equivalently, a function from centered worlds to truth-values). Such a content is 'poorer', less determinate, than a classical, full-fledged proposition. It is what, after Perry (1979/1993b : 42), I call a *relativized proposition*.

The debate over temporalism

A number of authors have agreed to relativize (the truth of) propositions to indices other than/richer than possible worlds. This move has been especially popular where time is concerned : in part because of Kaplan's argument, it has long been the majority view that tensed sentences express temporally-neutral propositions — propositions true at some times and false at others. The success of Priorean tense logics, based on that idea, has played a major role in making this the orthodoxy. But temporal relativism (or 'temporalism', for short) has met with considerable resistance in the recent philosophy of language, and the emerging consensus is that temporalism rests on a mistake. Since temporalism is supposed to be the most defensible form of relativism, rejecting temporalism amounts to rejecting relativism entirely.[3]

In Parts II and III I will present, and attempt to defuse, influential objections to temporalism : semantic objections (Part II) and epistemic objections (Part III).[4] This defense is part of my advocacy of a generalized, yet moderate, form of relativism which I am now going to describe in general terms. To do so, I start by considering what is perhaps the most basic objection to relativism : the Fregean objection, based upon the indeterminacy or incompleteness of the contents evaluated with respect to rich circumstances (e.g. centered worlds). As we shall see, there are two possible responses to that objection, and they correspond to two distinct brands of relativism : Moderate Relativism and Radical Relativism.

---

[3] By 'relativism', I mean the view that propositions can/should be relativized to indices other than/richer than possible worlds.

[4] A more detailed and systematic defense of temporalism against recent objections can be found in Berit Brogaard, *Transient Truths : An Essay in the Metaphysics of Propositions*, forthcoming.

Chapter 2

Radical vs. Moderate Relativism

The objection from incompleteness

Frege took fictional sentences to be endowed with content yet he considered them unevaluable. The author of a fictional statement does not attempt to characterize the actual world, he argued, so we are given a content without any circumstance of evaluation for it. One possible conclusion to draw from Frege's remarks on fiction is that, to get a truth-value, a content is not enough; we need to connect that content with the actual world, via the assertive force of the utterance, in virtue of which the content is presented as characterizing that world. Frege also held that, if a sentence lacks the force of a serious assertion (because the speaker does not attempt to characterize the actual world but is engaged in a different enterprise such as poetry), making the content of the sentence more complex by means of operators such as 'it is true that' or 'I assert that' will not change the situation. Whether or not an utterance is serious and characterizes the actual world is a pragmatic matter — a matter of 'force', not a matter of content.

From what I have just said, it follows that Frege accepted something like Duality. But he would have rejected Distribution, for Distribution entails that there are sentential contents that are not fully propositional. Such contents are unacceptable, from a Fregean point of view.

Frege rejected the very idea of a temporal proposition, i.e. a proposition that is true at some times and false at other times. Such a proposition is not a genuine proposition, he held, because it is not evaluable as true or false, or at least, it is not evaluable *unless* we are given a particular time. In the absence of a time specification, the alleged proposition is only 'true-at' certain times and 'false-at' others. It is, therefore, *semantically incomplete* by Frege's lights:

A thought is not true at one time and false at another, but it is either true or false, *tertium non datur*. The false appearance that a thought can be true at one

time and false at another arises from an incomplete expression. A complete
proposition or expression of a thought must also contain a time datum. (Frege
1967 : 338 ; quoted in Evans 1985 : 350)

As Evans points out, the problem of semantic incompleteness does not arise in the
modal case. Even if a thought is said to be 'true at' one world and 'false at' another,
as in modal logic, this does not prevent it from being true (or false) *tout court.* It is
true *tout court* iff it is true-at the actual world. But the 'thought' that it is hot cannot be
evaluated as true or false *tout court.* In the absence of a contextually supplied time it
can *only* be ascribed relative, 'truth-at'-conditions. Only a particular, dated utterance
of such a sentence can be endowed with genuine truth-conditions. What this shows
is that the time of utterance is part of the (complete) content of the utterance, or, in a
Fregean framework, part of the expression of such a content; hence it cannot be
considered external to content and treated like the world of evaluation. So the
objection goes.

The same objection evidently applies to the place-neutral content of 'It is
raining' : such a content is not complete, since the utterance cannot be evaluated
unless a place is contextually provided. The contextually provided location *is* an
ingredient of the (complete) content of the utterance, or, in a Fregean framework, it
contributes to the expression of such a content.

A first relativist response

There are two possible responses which a relativist can make, corresponding to two
forms of relativism. A radical relativist will insist that the partial contents we are led to
encompass if we accept Distribution *are* complete. Thus the Stoics posited '*lekta*' that
were « in many respects reminiscent of the 'propositions' that many modern
philosophers postulate as meanings of eternal assertoric sentences », save for the
fact that they were « temporally indefinite in the same way as occasion sentences »
(Hintikka 1973 : 70). Such *lekta* were thought by them to be complete, despite their
temporal neutrality. In his review of Mates 1953, which brought Stoic logic to the
forefront of attention, Geach wrote that for the Stoics, « though the truth-value of
'Dion is alive' changes at Dion's death, the sentence still expresses the same

complete meaning (*lekton*) » (Geach 1955 : 144).[5] This idea, which aroused Prior's interest, Evans later found incomprehensible and even incoherent. If the *lekton* is complete, Evans argued, it can be evaluated as correct or incorrect ; but if the *lekton* is temporally neutral, its evaluation as correct or incorrect will vary with time, hence it will not be evaluated as correct or incorrect once for all. Does this not entail that the *lekton cannot* be evaluated as correct or incorrect (tout court) after all ? Evans writes :

> To say that the sentence-type 'Socrates is sitting'… expresses a complete meaning seems to imply that… to know what assertion is being made by an utterance of a tensed sentence all you need to know is which tensed sentence was uttered ; you do not need further information to tie the sentence down to a particular time… It would follow that such an 'assertion' would not admit of a stable evaluation as correct or incorrect ; if we are to speak of correctness or incorrectness at all, we must say that the assertion is correct at some times and not at others (Evans 1985 : 349)

For Evans, this consequence (the instability of evaluation) is a *reductio* of the whole position. At this point, however, it is useful to consider MacFarlane's discussion of future contingents, for MacFarlane seems to bite Evans' bullet. MacFarlane (2003) argues that a sentence like 'There will be a sea-battle tomorrow'

---

[5] Geach's statement is alluded to in Prior 1967 : 17 and cited in Evans 1985 : 349. In an earlier review (Geach 1949), also quoted by Prior (1967 : 15) and Evans (1985 : 348), Geach had written that « for a scholastic, 'Socrates is sitting' is a complete proposition, *enuntiabile*, which is sometimes true, sometimes false ; not an incomplete expression requiring a further phrase like 'at time *t*' to make it into an assertion». That remark, later amplified by Geach in his review of Mates, played an important role in Prior's intellectual development : « Today this has perhaps become a commonplace of logical history, but in 1949 it was quite widely informative. It was certainly informative to myself ; I had taken it for granted that it was not only correct but also 'traditional' to think of propositions as incomplete, and not ready for accurate logical treatment, until all time-references had been so filled in that we had something that was either unalterably true or unalterably false. Geach's remark sent me to the sources. » (Prior 1967 : 15-16)

is neither true nor false when it is uttered (since the future is indeterminate) but turns out to be true or false, as the case may be, when it is evaluated the next day. So MacFarlane gives up the constraint that the evaluation of a thought as correct or incorrect must be temporally stable : the truth-value of an utterance may well depend upon the context of evaluation (e.g. the time at which it is evaluated), so an utterance or thought that is evaluated in a certain way at a certain time may be evaluated differently at a different time. In such a framework, reminiscent of Aristotle, we could maintain that the tensed sentence 'Dion is alive' expresses a complete content, and is (therefore) evaluable (at any given time), since we reject the constraint that the evaluation process itself must be 'eternal', hence stable, rather than context-sensitive and unstable.[6]

A second relativist response

Whatever we think of this Aristotelian line of argument, I will not be concerned with the radical forms of Relativism in this book, but only with a moderate form which I myself advocate. In response to the Fregean objection, a moderate relativist will *concede* that the complete content of the utterance/thought 'Dion is alive' involves more than the temporally neutral *lekton* it expresses ; it additionally involves the time of utterance, which is tacitly referred to and against which the utterance is meant to be evaluated. Distribution can be construed as saying that the complete content, *in*

---

[6] Note, however, that this is not what MacFarlane himself would say. MacFarlane's brand of Radical Relativism consists in making room for a new form of context-sensitivity : sensitivity to the context of evaluation and not (or not merely) to the context of utterance. On MacFarlane's view, some expressions are 'assessment sensitive', and others are not — just as some expressions are 'utterance sensitive', and others are not. In this regard, future contingents are a *special* case. (Evans himself seems to accept that there is something special about future contingents, and for that reason, he says, he confines his discussion to sentences in the past. See Evans 1985 : 350, fn. 9.) As far as 'Dion is alive' is concerned, MacFarlane holds that its truth depends on the time of utterance, as Frege points out, but *not* on the time of assessment. This shows that one may be a radical relativist with respect to some sentences — those whose truth-value is assumed to depend upon the context of assessment and whose evaluation is therefore unstable — and not with respect to others.

the sense of Frege and Evans, distributes over the two components which Duality posits, namely the circumstance of evaluation (which may include more than a world) and the content to be evaluated, in the narrow sense of content. Once it is admitted that we need these two components, we can tolerate contents that are not 'semantically complete' in Frege's sense, i.e. endowed with absolute truth-conditions. We can, because the circumstance is there which enables the content to be suitably completed. Thus the content of tensed sentences is semantically incomplete, yet the circumstance (the time) relative to which such a sentence is evaluated is sufficient to complete it. It follows that we must distinguish *two* levels of content. The content we evaluate with respect to the circumstance is the content in the narrow sense; it may, but need not be, semantically complete by Frege's lights. What is semantically complete in any case is the content in the broad sense. It consists of the (narrow) content *and* the circumstance with respect to which that content is meant to be evaluated. Distribution only induces us to *analyse* the complete content of an utterance into two components, corresponding to those distinguished in Duality.

Chapter 3
Two levels of content

Aristotle vs Frege

In his interesting article on 'Time, Truth and Knowledge in Aristotle and Other Greek Philosophers', Hintikka presents the Aristotelian view as follows :

> It is obvious that the sentence, 'It is raining', as uttered by me today, is made true or false by a set of facts different from those that verified or falsified my utterance yesterday, 'It is raining'. But it is very natural to say that in some sense the state of mind or attitude toward my environment that is expressed by the two utterances is the same. The facts to which yesterday's utterance refers are referred to today by the sentence, 'It was raining yesterday'. But the 'state of mind' that this utterance appears to express seems to be entirely different from that expressed by yesterday's present-tense utterance, 'It is raining'. (…) Hence the idea that spoken words are symbols for unspoken thoughts encourages the idea that one and the same temporally indefinite form of words expresses one and the same belief or opinon at the different times when it is uttered. (Hintikka 1973 : 85)

For Aristotle as well as for the Stoics, the thought expressed by 'It is raining' stays constant on distinct occurrences. The truth-value changes, but that does not show that the thought in question is somehow incomplete. The truth-value changes because *the facts* change. The truth-value of a thought depends upon the correspondence (or lack of correspondence) between that thought and the facts. Since the facts change (from rain to nonrain) as time passes, the change in truth-value does not entail that there is a change also in the expressed content.

From a modern, Fregean point of view, this way of looking at the matter is unacceptable. A complete thought content can only be true or false, *tertium non datur*. Now the thought that it is raining is true (when entertained on a rainy day) *and*

false (when entertained on a sunny day) ; it follows that that 'thought' in question is not a genuine thought — a complete content — by Fregean standards. The shift in truth-value is sufficient to show that the thought expressed (in the Fregean sense of 'thought') changes from one occurrence to the next.[7]

Aristotle did not have the modern, Fregean notion of thought content at his disposal. Hintikka argues that that notion would have been most useful to him, given his doctrine that only substances can take contrary attributes at different times and nevertheless remain numerically one and the same (*Categories* 5. 4ᵃ 10 ff). The changing truth-value of the thought that it is raining is a prima facie counterexample, which Aristotle could have disposed of by appealing to the Fregean notion, had such a notion been available to him. Instead, Aristotle ruled out the counterexample by means of what Hintikka takes to be 'a rather unsatisfactory manœuvre'. « The fact that the modern view was not even considered by the author of the *Categories*, although it would have served his purpose perfectly, illustrates the hold of the contrary view on him » (Hintikka 1973 : 68).

A synthesis

Contrary to Aristotle, we do have the Fregean notion of (complete) 'thought' at our disposal. But that does not mean that we should reject anything in the quotation from Hintikka with which this chapter began. That quotation emphasizes the naturalness of the Aristotelian notion of thought content. It is, indeed, natural and appealing, but acknowledging this is compatible with the adoption of a broadly Fregean framework. That is so because the Fregean and the Aristotelian notions of thought content need not be seen as competing. It is possible to construe them as complementary notions, and such conciliation results in the framework I call 'Moderate Relativism'.

To a modern reader, familiar with the Fregean notion of content, the Hintikka quotation suggests that the complete content of an utterance (that which determines

---

[7] This notion of content as eternal is not only Frege's ; it is also Russell's. See Prior 1957 : 110 for a long quotation from Russell's review of MacColl (a reference which Prior found in Geach's review of Mates 1953). This is part of an appendix to Prior's *Time and Modality,* entitled 'Tenses and Truth in the History of Logic', and containing an interesting overview of the transition to the modern conception of proposition (Prior 1957 : 104-22).

its truth-conditions, i.e. the 'facts' which make it true or false) involves two factors : the (Aristotelian) thought that is expressed, and the time at which it is expressed. The sentence 'It is raining' expresses the same thought whenever it is uttered, and that thought is evaluated with respect to the time of utterance. Since the latter changes, the truth-value is liable to change even though, in the narrow sense of content, the content is the same. The truth-*conditions* also change : An utterance of 'It is raining' at t is true iff the thought expressed by the sentence is true at t ; an utterance of the same sentence at t' is true iff the same thought is true *at t'*. On this view the complete content of two successive utterances of 'It is raining' need not be the same, since the speaker does not merely express a certain content, but also tacitly refers to a certain time (the time of utterance) as relevant for the evaluation of that content. The complete content, individuated *à la* Frege, corresponds to the utterance's truth-conditions which, according to Hintikka, depend upon an external factor, namely the time at which the utterance is made or the thought entertained.

Dummett on Prior

In a similar vein, Dummett attempts to make sense of Prior's position (in response to Evans 1985) by distinguishing two levels. He points out that temporal propositions are, for Prior, the contents of *sentence-types*. The content of a sentence-type is a function from times to truth-values, hence a sentence-type only has relative truth-conditions: it is true at some times and false at other times. This does not prevent us from introducing a notion of absolute truth, by shifting to the level of *utterance* content. According to Dummett, when a sentence is uttered the function which is its content is applied to some contextually provided time (typically, the time of utterance). The time in question serves as circumstance of evaluation for the utterance: the utterance is true *tout court* iff the sentence is 'true-at' the contextually provided time. As Dummett points out,

> The variable truth-value and the absolute truth-value attach to different things; it is the type sentence that is true at one time, false at another, but the utterance that is true or false simpliciter (Dummett forthcoming : 44)

Since there are two distinct levels, corresponding to the sentence-type and the utterance, there is no harm in taking the utterance to possess a 'content' also (content$_u$), distinct from that of the sentence (content$_s$). For example, we can treat the utterance as expressing a structured proposition consisting of (i) the contextually provided time as subject, and (ii) the content of the sentence-type, predicated of that time. But if we do so, we must acknowledge the unarticulated nature of the 'subject' in the content$_u$ of tensed utterances. As Prior says, "tensed propositions are understood as directly or indirectly characterising the *un*mentioned time of utterance" (Prior 1977 : 30). Hence there is a trade-off: if we want to restrict ourselves to what is linguistically articulated, we must focus on the content$_s$, which is 'semantically incomplete' by Frege's lights — it corresponds to the content of a predicate rather than to that of a complete sentence in a logically perfect language. If, following Frege, we want to focus on the complete content of the utterance, that which makes it truth-evaluable in absolute terms, we must acknowledge the role played in that content (content$_u$) by unarticulated constituents corresponding to the circumstances in which the content$_s$ is evaluated.

Barwise and 'Austinian propositions'

Another author who ought to be mentioned in connection with Moderate Relativism is Jon Barwise. Barwise also put forward a semantic theory with two levels of content : the 'infon' or 'state of affairs', which is the content in the narrow sense, and the richer 'Austinian proposition'. « Whenever we make a claim », Barwise and Etchemendy write, « it is a claim about some situation or other » (1987 : 176), so « all propositions contain an additional contextually determined feature, namely, the situation they are about » (1987 : 29). The infon is the content to be evaluated with respect to a given situation — the claim we make regarding that situation — and the Austinian proposition is the proposition to the effect that that situation supports that claim. In what follows I will use the notion of Austinian proposition, corresponding to the complete content of an utterance/thought. But I will use the Stoic term '*lekton*', rather than Barwise's theoretically-loaded term 'infon', to refer to the content in the narrow

sense.[8] So 'It is raining' expresses a constant *lekton* whenever and wherever it is used, a content that can be modeled as a function from situations to truth-values or as a set of situations (viz. the set {s : it is raining in s}); but the complete content of an utterance of 'It is raining' is the Austinian proposition that a certain situation (that which the utterance/thought 'concerns') fits that *lekton*, i.e., belongs to the set of situations in question.

---

[8] Evans also has coined a term for that entity. He calls it the 'Stoic-proposition'. (Salmon 1986 : 24-27 calls it a 'proposition matrix'.) As for 'infon', the term is originally due to Keith Devlin (see Devlin 1991). In the introduction to his collection of papers *The Situation in Logic*, Barwise writes : « The terms 'state of affairs' and 'infon' are used interchangeably in different chapters of this book. Somewhere during the decade we decided that the former had too many unwelcome connotations » (Barwise 1989b : xv). Barwise and Etchemendy (1987) still use 'state of affairs'.

Chapter 4
Branch points for Moderate Relativism

Is the lekton *context-dependent ?*

In Dummett's framework the relativized content — which I call the *lekton* — is the content of the sentence-type. That content is admittedly incomplete : it is like the content of a predicate. To get a complete content we need a situation, which the utterance context provides and the *lekton* can be predicated of. Change the situation, you change the complete content of the utterance, even though the content of the sentence (the *lekton*) remains constant.

Barwise's view is very similar, but there is a significant difference. The infon, according to Barwise, is the content of the sentence *with respect to context* : if the sentence contains indexicals, the contextual values of the indexicals contribute to the infon. On this interpretation, the *lekton* itself is context-dependent, and we have three levels rather than merely two. The three levels are : the meaning of the sentence-type, the *lekton*, and the complete Austinian proposition. On this version of Moderate Relativism, which I will elaborate in this book, the context comes into the picture twice : first, it provides values for the indexicals, which values contribute to the *lekton* ; second, it determines the situation against which the *lekton* is to be evaluated. The complete content of the utterance involves the *lekton* together with the situation of evaluation.

So there are two possible interpretations for the notion of *lekton* within the moderate relativist framework, depending on the option we choose with respect to the issue I have just raised. Is the *lekton* context-dependent (Barwise), or can it be equated with the meaning of the sentence-type (Dummett)? The first option corresponds to what MacFarlane calls 'Non-indexical Contextualism' (MacFarlane forthcoming). It a view famously held by Kaplan, since Kaplanian 'contents' do not determine a classical proposition unless a circumstance (involving a time and, possibly, a place in addition to a world) is contextually provided. The second option corresponds to Lewis's position, as we shall see in Part IV, when we discuss the

proper treatment of indexicality, and again in Parts X and XI, when we discuss the analysis of *de se* thought.

As far as I am concerned, I will follow Barwise and Kaplan and construe the *lekton* as the content of the sentence in context, so that an indexical sentence will express different *lekta* in different contexts. My main reason for taking the *lekton* to be an extra level of content, intermediate between the meaning of the sentence and the complete content, is that this enables us to represent what the sentence (or possibly the thought) explicitly articulates (in a possibly indexical manner). Again, 'It is raining here' says something different from what 'It is raining' says, even in a context in which they are both true iff it is raining at the place of utterance. The difference lies in the fact that the place in question is (indexically) articulated in the former case while it is left unarticulated in the latter. It follows that the *lekton* differs, even though the truth-conditions are the same.

What do we do when the lekton *is a classical proposition?*

The second branch point concerns the special case in which the content of the sentence is not semantically incomplete, but is already a fully-fledged, classical proposition. In such a case the *lekton* determines a function from possible-worlds to truth-values, hence the only thing we need to determine a truth-value is a possible world. No further relativization is needed.

It seems that, with sentences whose content is not semantically incomplete, there is no need to invoke a double layer of content. The content of the sentence, insofar as it has an absolute truth-value, is the only thing we need. Or, to put it in slightly different terms : in such cases the *lekton is* the complete content. This position defines one version of Moderate Relativism, namely the weak version ('WMR', for 'weak moderate relativism'). But there is another, strong version, which has been argued for by Barwise and which I also advocate.

Before proceeding, we should note that this issue is orthogonal to the first issue I raised (the context-dependence of the *lekton*). If we are context-independentists, the special case we are talking about is the case in which the content of the sentence-*type* is a classical proposition, i.e. the case in which the uttered sentence is eternal. If, like Barwise and Kaplan, we take the *lekton* to be context-dependent, then the special case at issue is the case in which the content of

the *sentence-in-context* is a complete proposition. Suppose the speaker says 'It is raining here and now', without leaving anything unarticulated (save the world of evaluation). Both the time and the place are explicitly articulated, hence they are both part of the *lekton*. It follows that we do not need a rich circumstance to evaluate that content : we only need a possible world. Here the sentence is not eternal, but it is fully articulate, so that no further relativization is needed.

Whether we construe the *lekton* as context-dependent or not, the special case in which it is a classical proposition can be handled in two ways. Weak Moderate Relativism says that, in such a case, the *lekton* is the utterance's complete content ; so the distinction between the two levels of content only holds in the cases in which the *lekton* is not classical. Strong Moderate Relativism ('SMR') rejects that conclusion. According to SMR, the content of a sentence (*whatever* the sentence) is a function from situations to truth-values. Hence the relativity of truth, construed as a property of sentences: the same sentence may be true relative to a situation and false relative to another one. That is so *even if the sentence itself is not semantically incomplete*.

Strong Moderate Contextualism

Even when the sentence *is* truth-evaluable in the absolute sense — when it is 'semantically complete' by Frege's lights — SMR says there is a principled distinction between the content of the sentence (the *lekton*) and the content of the utterance (the Austinian proposition). In such a case, the *lekton* will be a 'classical' proposition (a function from possible worlds to truth-values), but the Austinian proposition will still contain a situation in addition to that proposition. What the utterance 'says' is that *the situation in question supports the proposition in question.* It follows that two distinct evaluations are possible, in such cases. We can evaluate the sentence itself (i.e. evaluate the proposition with respect to the actual world), or we can evaluate the utterance, that is, evaluate the proposition *with respect to the situation figuring in the Austinian proposition*.

To illustrate this point I usually quote my favourite example, from Barwise and Etchemendy (1987 : 29, 121). Commenting upon a poker game I am watching, I say: 'Claire has a good hand now'. What I say is true, iff Claire has a good hand in the poker game I am watching at the moment of utterance. But suppose I made a

mistake and Claire is not among the players in that game. Suppose further that, by coincidence, she happens to be playing bridge in some other part of town and has a good hand there. Still, my utterance is not intuitively true, because the situation it concerns (the poker game I am watching) is not one in which Claire has a good hand at the time of utterance. But we can say that the *sentence* (or the sentence-in-context) is true: for it says that Claire has a good hand at the time of utterance, and Claire *has* a good hand (somewhere) at the time of utterance. The unarticulated constituent which distinguishes the *lekton* from the Austinian proposition makes all the difference here, and it accounts for our intuitive classification of the utterance as non-true.

This sort of approach can easily be extended to deal with standard problems such as that of quantifier domain restriction. It is natural to hold that 'All *F*s are *G*' expresses a proposition that is true (in a world, at a time) if and only if the *F*s are all *G* (in that world, at that time). Thus 'All students are French' expresses the proposition that all students are French. Many theorists feel compelled to give up this natural view, and claim that the sentence is semantically incomplete or covertly indexical, so that it expresses no proposition (independent of context).[9] They say so because they are impressed by the fact that the truth-conditions of an utterance of that sentence typically involve a contextually restricted domain of quantification. In the SMR framework, however, we can stick to the simple and straightforward view regarding the proposition expressed by 'All *F*s are *G*', while fully acknowledging contextual domain restriction. The two layers of content enable us to do just that. The sentence is said to express a proposition that is evaluable with respect to an arbitrary world (or, perhaps, an arbitrary world-time pair) — the proposition that all students are French — but that proposition can *also* be evaluated with respect to the specific situation that features in the Austinian proposition. That is what happens when we evaluate an utterance of this sentence, instead of evaluating the sentence itself.[10]

---

[9] See e.g. Stanley and Szabo 2000.

[10] A well-known difficulty for the situation-theoretic approach to contextual domain restriction comes from the fact that distinct quantifiers in a single sentence may involve distinct restrictions. The answer to that difficulty consists in associating sub-sentential expressions with (local) circumstances of evaluation. See e.g. Recanati 1996.

Part Two

The debate over temporalism (1) :

Do we need temporal propositions ?

Chapter 5

Modal vs extensional treatments of tense

It is common nowadays to reject the idea that tenses should be represented as analogous to the sentential operators of modal logic. This idea, on which Arthur Prior built the new discipline of tense logic in the middle of the twentieth century, had been widely accepted among philosophers for almost twenty years, but the development of formal semantics as an autonomous discipline within linguistics, and the widespread adoption of a non-modal framework for dealing with tenses by practitioners of that discipline, led to a gradual abandonment of the idea.

Now, if the tenses are *not* operators, then the argument that we need time-neutral propositions for such operators to operate on collapses. The abandonment of the tense-logical approach by practitioners of formal semantics has thus been used to argue against temporal propositions. A recent instance of that line of argument can be found in King's paper 'Tense, Modality, and Semantic Values' (King 2003).

In Part II I attempt to show that Temporal Relativism (and Moderate Relativism in general) is not threatened by arguments along those lines. I start by presenting the reasons that have been adduced for giving up the modal framework in the theory of tense.

Partee's objection to Prior

A first objection to the Priorean treatment of tense is due to Barbara Partee (1973). She argues that a certain class of examples, in which reference is made to a specific time (indicated by the tense of the verb), speak against a treatment in terms of operators. In her most famous example, 'I did not turn off the stove', the past tense is understood *deictically* : the speaker refers to a specific time in the past. As Partee writes,

> When uttered, for instance, halfway down the turnpike, such a sentence
> clearly does not mean either that there exists some time in the past at which I

did not turn off the stove or that there exists no time in the past at which I turned off the stove. The sentence clearly refers to a particular time. (Partee 1973/2004 : 51)

Tense-logical operators, Partee says, cannot capture the referential nature of these uses of tenses. The 'past' operator of tense-logic presents the proposition it operates on as true with respect to *some* time in the past, but does not refer to a particular time. In contrast, she points out, there is a striking parallel between tenses and pronouns, a parallel that justifies treating tenses as variables rather than operators. Like pronouns, tenses have not only deictic, but also anaphoric and bound uses. The existence of all three uses suggests that it would be more fruitful to exploit the analogy with pronouns and to handle tense in terms of reference to, and quantification over, times, rather than in terms of temporal modalities.

Interesting though they are, Partee's remarks are not sufficient to dispose of the tense-logical approach. Even if we hold that the past tense functions like an operator which presents its operand as true with respect to *some* time in the past, we can easily account for prima facie deictic uses. The deictic reading arguably results from restricting the implicit domain of quantification in a manner that mimics singular reference. Thus 'I didn't turn off the stove' can be analysed as saying that, in the contextually relevant period (that which immediately follows my last use of the stove), there is no time at which 'I turn off of the stove' is true. On this analysis what the sentence encodes is the general proposition that there is no time at which I turn off of the stove, and the context provides a temporal situation for that statement to be evaluated against (for a recent proposal along those lines, see Lasersohn 1999 : 537 fn).[11]

Note that tenses, unlike pronouns, admit of 'existential' uses. The future tense sentence 'I will go to China' may be understood in two ways : either the speaker is referring to some specific future time, made salient in the context, and she says that she will go to China at that time ; or she is merely saying that *there is* a future time at

---

[11] Partee herself mentioned a similar analysis, giving narrow scope to the negation : « It might be possible to… contend that the sentence asserts only that there is some time in the past at which I did not turn off the stove, with the narrowing down to relevant times explanable by conversational principles. » (Partee 1973/2004 : 57 fn.)

which she will go to China.[12] Such an existential reading is not available with pronouns : 'He is bald' cannot mean that some male or other is bald, even within a clearly identified group of people. This casts some doubt on the 'tense as pronoun' analysis.

To sum up, if we say that the tenses are always (implicitly) quantificational, because they are fundamentally operators, we can still account for the 'deictic' uses in terms of contextual domain restriction ; so we can give a unified analysis of existential and deictic uses. (As for anaphoric uses, they can be handled in the same way as the latter.) But if tenses are to be treated as referential variables (as Partee 1973 suggests), the variables in question must be different from pronominal variables : they must be *optional* variables. An optional variable, when unbound, may be contextually assigned a specific value *or* undergo existential closure. Whether there are such variables is unclear, and can be disputed (see Recanati forthcoming a).

'Referential' operators in the tense-logical framework

Be that as it may, let us assume that Partee is right and that the tenses have a genuinely referring use, irreducible to their existential use. This fact, by itself, does not show that the tenses cannot be analysed as tense-logical operators. To be sure, the *standard* temporal operators — the 'P' and 'F' operators of Prior's tense logic — are (implicitly) quantificational rather than referential ; but nothing prevents the introduction of 'referential' operators in the tense-logical framework. We find seeds of that idea in Prior's own work (see e.g. Prior 1967, chapter 6, and Prior 2003 :183,

---

[12] The past tense sentence 'I did not turn off the stove' can also be given such an interpretation, rather than a deictic interpretation. One week after the event reported by her utterance 'I did not turn off the stove', Partee might have told the story as follows : 'Last week, I did not turn off the stove and I risked a serious accident'. Even though the sentence contains a reference to the time-span within which the event talked about is said to occur ('last week'), there is no deictic reference to the time of the event itself, which remains indefinite.

227-9),[13] and recently tense-logicians have developed extensions of Priorean tense logic in which arbitrary reference to times is possible (see e.g. Blackburn 1994). In such a hybrid system one can have a tense-logical operator '@*t*' such that, when it is prefixed to a sentence, the result is true iff the embedded sentence is true at time *t* (Blackburn 2000 : 347-8).[14] And of course, early in the development of tense logic, Hans Kamp had introduced indexical operators such as 'N' (for 'now') which refer to specific times (the time of utterance, in the case of 'N'). Among the expressions which function like 'now' and can be modeled in the same sort of way, Kamp mentioned 'then', which is the counterpart of a pronoun in the temporal domain (Kamp 1971 : 230, Vlach 1973). So the possibility of introducing 'referential' operators into tense-logic alongside the traditional 'quantificational' operators was already taken on board even before Partee raised her objection to the tense-logical approach. It follows that that objection is not an objection to the tense-logical approach per se, but only to a particular version of that approach.

Explicit quantification over times

That is not to deny that we can account for the tenses by construing them as pronominal devices in an extensional framework, as Partee 1973 suggested. In his pioneering paper 'Formal Properties of 'Now'', Kamp himself briefly mentioned (in a footnote) « the possibility of [using] explicit quantification over moments. Such a symbolization… would certainly be possible ; and it would even make the operators P and F superfluous » (Kamp 1971 : 231). It is interesting in retrospect to look at the two objections which Kamp then levelled against such an alternative approach —

---

[13] Evans notes that Prior's 'metric tense logic' makes it possible to represent referential uses of the past tense : « [Partee] wrongly thinks that [the existence of such uses] rules out the recognition of tense operators in English, forgetting Prior's metric tense operators » (Evans 1985 : 361 fn.)

[14] In this system '*t*' is a propositional atom of a special sort, a 'nominal', true at exactly one time in any model. Though they are propositional atoms, nominals have a referential function : they name (or at least, identify) the only state they are true at. The '@'-operator allows one to shift to the state thus named and to evaluate a formula there. For a technical introduction to hybrid logic, see Blackburn, de Rijke and Venema (2001 : 434-45).

precisely the approach that has gained wide acceptance among linguists and philosophers recently. As we shall see, these objections are far from convincing.

About the extensional approach with explicit quantification over moments, Kamp wrote :

> Such symbolizations 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. Moreover, one can object to symbolizations involving quantification over such abstract objects as moments, if these objects are not explicitly mentioned in the sentences that are to be symbolized. (Kamp 1971 : 231 fn.)

The second objection is a Quinean sort of objection : one should not quantify over, hence ontologically commit oneself to, dubious entities such as moments or possible worlds. But from the point of view of Prior, a quantified statement commits us to no more and no less than its substitution instances do. So we can generalize from 'I hurt him by treading on his toe' to 'I hurt him somehow' or 'I hurt him in some way'. The existentially quantified statement we end up with does not commit us to the existence of 'ways' *anymore than* the original sentence 'I hurt him by treading on his toe' does (Prior 1971 : 37-39). Given Prior's 'non-objectual' understanding of quantification,[15] the Quinean argument that one should not quantify over objects that one may not be prepared to countenance in one's ontology has no force.

What about the first argument ? Prior insisted on treating the tenses as modifiers akin to adverbs rather than as pronominal elements. The logical form of tensed sentences analysed *à la* Prior is simpler than the logical form that results from explicitly quantifying over moments, and this matches the corresponding simplicity of the natural language sentences. The formulas one gets when one explicitly quantifies over moments are simply too long ! However this argument has lost much of its weight. Tense theorists have developed frameworks with explicit quantification over times in the object-language, and such frameworks turn out to offer elegant and perspicuous ways of representing complex facts about natural-language tenses (for example, sequence-of-tense phenomena, which King uses as example in his 2003

---

[15] For a lucid presentation of Prior's view, see MacBride forthcoming.

paper)[16] — facts which the tense-logical approach can handle only by introducing extra complexities which offset the initially appealing simplicity of the approach. Thus I said that it is possible to account for referential and anaphoric uses of tenses within the tense-logical framework. To do so, however, one arguably has to introduce a series of new operators

> [whose definition] tends to become very complicated and [whose] number tends to grow with every new phenomenon that cannot be expressed with the previous operators. For that reason, more and more people take a language with quantification over times as basis, and just define the operators they need in them… For the simple cases this has the disadvantage of length… For the longer cases that has the advantage of readibility and keeping track of the interpretation (Landman 1991 : 135)

To be sure, the tense-logical formulas are closer to *surface grammar* than their long-winded extensional counterparts — that is, presumably, what Kamp meant when he said that the latter 'depart from the actual form of the original sentences which they represent' — but contemporary semanticists mostly care about *another* level of syntactic representation, namely, that which is the input to semantic interpretation (the so-called 'Logical Form', distinct from surface grammar).

I conclude that there are, indeed, two possible approaches to tense : the tense-logical approach whose operators implicitly quantify over (and possibly refer to) times but do not do so explicitly ; and the extensional approach with explicit quantification over (and reference to) times in the object-language. Far from being the only correct way of representing natural language tenses, the tense-logical approach tends to be disfavoured as less flexible and perspicuous than the other approach. It follows that one may, but certainly *need not*, construe the tenses as operators. So Kaplan's argument that we need temporal propositions for such operators to operate on weighs very little. Or rather, it has weight only for someone who is antecedently committed to the correctness of the tense-logical approach as opposed to the other approach. There still are such people, but they are a minority in

---

[16] Higginbotham also adduces these facts in support of his 'reflexivist' analysis of tensed sentences, to be criticized below (Chapters 22-26).

the present theoretical landscape. That is the gist of King's argument against temporal propositions.

Chapter 6
What is at stake?

We have just seen that there are (at least) two theoretical options for the treatment of tense. Basically, one is free to adopt one framework or the other. As Ogihara writes,

> [The] choice of logical language should not be taken as an important theoretical decision… The only important issue is whether the language has enough tools to describe the target construction in natural language. (Ogihara 1996 : 28, quoted in King 2003 : 218)

King disapproves of this 'pragmatic, almost instrumentalist' attitude. He holds that there are *substantive reasons* for giving up the modal approach. I disagree. In this chapter and the next one, I will argue that, to the extent that there are deep theoretical reasons at stake, they favour the approach in terms of operators.

The expressive power argument

For a long time there has been a theoretical argument around, to the effect that the extensional approach is to be preferred, on grounds of expressive power. The argument runs as follows. Using the technique of 'Standard Translation' modal languages may be shown to be fragments of standard first-order languages (van Benthem 1985). If natural language tenses can be dealt with within a modal framework, as Prior suggested, it would seem to follow that the full power of classical logic is not needed. Conversely, if the full expressive power of first-order logic is needed to deal with natural language tenses, it would seem to follow that they cannot be dealt with in a modal framework. Starting with these premisses, the argument is straightforward. It has been established that the full expressive power of first-order logic *is* needed to deal with natural language tenses. (See e.g. van Benthem 1977, Cresswell 1990, chapter 2-4.) Does it not follow that natural language tenses cannot be dealt with in a modal framework, and that we need to go extensional ?

No, it does not ! The conclusion does not follow because, as we have seen, the modal language can be *enriched* with new operators, such as Kamp's 'N' operator (corresponding to 'now') or Vlach's 'K' operator (corresponding to 'then'). Such an enrichment can proceed until the modal language one works with has itself acquired the full power of standard first-order languages. Quine has demonstrated that this is possible, since quantification can be expressed without variables, by using suitable operators (Quine 1960). Thanks to those operators, first-order languages themselves can be translated into modal languages (Kuhn 1980). Exploiting the Quine-Kuhn techniques, one can therefore provide modal languages in general (and tense-logical languages in particular) with the full expressive power of standard first-order logic (Cresswell 1996, chapter 1). So the formal argument that we need explicit quantification over times on grounds of expressive power is a nonstarter.

King's argument

King is well-aware of all this, and the expressive-power argument is not what he appeals to when he claims that the tense-logical approach is no longer a live option. He argues as follows :

> My claim that English tenses cannot be viewed as operators cannot be based on the claim that to treat them in this way would be to not capture the expressive power of English… Rather, the claim is that treating tenses as e.g. involving quantification over times (and expressing relations between times) rather than index shifting sentence operators (i) allows for a simpler, more elegant, less ad hoc treatment of tenses and temporal expressions than does an operator treatment ; and (ii) allows for a more plausible account of the relation between the surface structures of English sentences and the syntactic representations of those sentences at the level of syntax that is the input to semantics… This, in turn, explains the fact mentioned above : that virtually all current researchers trying to give a treatment of the complex temporal data in natural languages eschew an operator approach to tenses in favor of treating tenses as something like quantifying over, referring to and/or expressing relations between times. (King 2003 : 221)

King overstates his case, however. He invokes a majority of semanticists who indeed prefer the extensional approach to the modal approach on overall grounds of elegance and simplicity, and concludes that « if the complex temporal facts present in natural language are most readily and easily represented by viewing tenses as involving explicit quantification over time and as expressing relations between times, that is a good reason for thinking that tenses really work this way » (King 2003 : 218). King therefore dismisses Ogihara's pragmatic attitude by appealing to the very preferences which Ogihara and his colleagues manifest. But such preferences are perfectly compatible with, indeed can hardly be dissociated from, the pragmatic attitude in question. Everything being equal, one naturally prefers the framework that is most convenient given one's other commitments in syntax and semantics. Natural and important though they are, such preferences remain superficial and Ogihara's pragmatic attitude toward them is perfectly appropriate.[17] If there were *deep conceptual reasons* for preferring the modal framework, they would undoubtedly override such pragmatic preferences. So the question we must ask is : Are there, or are there not, deep conceptual reasons for treating the tenses as operators ?

The internal perspective

I think there is at least one, and it has been concisely stated by Patrick Blackburn (one of those who still work within Prior's tense-logical framework). Blackburn writes :

> Prior insisted on the primacy of the *internal* view of time. This view situates the speaker firmly inside the temporal flow : the speech-time centred ordering of *past*, *present* and *future*, rather than the absolute *earlier-than*/*later-than* relation, is considered central to the analysis of tensed talk. This is natural : we live *in* time, and the internal perspective is imprinted on natural language in many ways.
>
> Prior's insistence on the internal perspective led him to develop his temporal calculi as modal logics, not classical logics. Classical logic, with its explicit variables and binding, offers a God's-eye-view of temporal structure ;

---

[17] Note that there are also pragmatic reasons, such as compactness or computational tractability, for preferring the modal approach. See Blackburn (1994, 2000).

modal logic, on the other hand, uses operators to quantify over structure 'from the inside'. (Blackburn 1994 : 83)

Following Blackburn, I will argue that the main reason one has for treating the tenses as operators is that, in this way, one captures the internal perspective on time which characterizes our thought and talk. This is independent of any preference one may have for or against the 'A-theory' of time and the metaphysical view Prior called 'presentism'. One may be a 'B-theorist' and hold that, metaphysically, the earlier-than/later-than relation is all there is to time, while still accepting that, from the cognitive point of view, the internal perspective is primary. Such a view has been argued for by Hugh Mellor (1998). Whether one follows Prior or Mellor in the metaphysics of time, in both cases one may accept, as a deep constraint on our theorizing, the need to capture the primacy of the internal perspective in dealing with the representation of time in language and thought.

By way of introduction to this chapter, let me give another quote from Blackburn, this time from his 'Hybrid Logic Manifesto' :

> Modal languages talk about relational structures in a special way : they take an *internal* and *local* perspective on relational structure. When we evaluate a modal formula, we place it *inside* the model, at some particular state *w* (the *current state*). The satisfaction clauses (and in particular, the clause for the modalities) allow us to scan other states for information — but we're only allowed to scan states reachable from the current state. The reader should think of a modal formula as a little automaton, placed at some point on a graph, whose task is to explore the graph by visiting accessible states. (Blackburn 2000 : 341-2)

What matters to us — and what I will focus on in this chapter — is the dissymmetry between the current state and the other accessible states, to which the modal operators take us. The modal operators shift the point of evaluation from the current state to some other accessible state. The dissymmetry consists in the fact that the current state is antecedently *given* and does not have to be represented, in contrast to the other states which have to be explictly introduced for the shift to take place.

Unarticulatedness

That the point of evaluation for a sentence is not, or at least does not have to be, represented in that sentence is one aspect of the 'Distribution' idea. Let us first consider the simple case in which the point of evaluation is a possible world. Clearly, the possible worlds relative to which propositions are evaluated are not themselves represented in, or an aspect of, the propositions in question. Thus 'I am French' is true, with respect to a world *w*, iff I am French in *w* ; but the sentence 'I am French'

only talks about me and the property of being French. The world of evaluation is not a constituent of the content to be evaluated.

To be sure, one can bring the world into the content by making the statement more complex. The complex sentence 'Possibly, I am French' (or 'I might be French') tells us that *in some possible world* I am French. The modal statement I make by uttering that sentence is about possible worlds, not merely about me and the property of being French.[18] In hybrid logic, one can even make statements 'referring' to specific possible worlds, just as it is possible to make statements referring to specific times. The hybrid-logic formula '$@_w p$' tells us that p is true at world w. The modal statement 'Actually, I am French' itself can be seen as talking about the actual world as well as about myself and my nationality. But the worlds that are thus introduced into the content of the complex statement (via modal operators such as 'possibly', 'actually' or '$@_w$') are used in evaluating the *simple* statement that is embedded within the modal statement. The modal statement itself is evaluated, and it shares with the simple statement the property that *the worlds with respect to which it is evaluated are not themselves represented in the statement under evaluation*.

To appreciate the unarticulated character of the circumstance of evaluation in the modal framework, it is worth looking at what happens when, using the techniques of Standard Translation, we transpose a modal statement into first-order logic, by explicitly quantifying over possible worlds. Simplifying somewhat, the translation runs as follows : 'Necessarily p' becomes '$\forall w\, p(w)$', 'Possibly p' becomes '$\exists w\, p(w)$', etc.[19] All complete sentences are transformed into predicates (of worlds). A simple categorical statement such as 'Rain is wet' will be represented as '$p(w)$', where 'p' is the proposition that rain is wet transformed into a predicate of worlds, and 'w' is a free variable to which the actual world is contextually assigned as default value.

The big difference between the modal statement and its standard translation is that, in the extensional framework, the circumstance of evaluation (the world) becomes a constituent of content. The contrast between content and circumstance is lost. Now that contrast makes a lot of sense. To evaluate a sentence, we determine whether the state of affairs it describes obtains in some 'reality' which serves as

---

[18] Prior would deny that it is 'about' possible worlds in the same sense in which it is about me and the property of being French. For the sake of the argument, I ignore this complication.

[19] To keep things simple I skip accessibility relations (or rather, I leave them implicit).

circumstance of evaluation. But that reality — the actual world, say — is not itself, or at least doesn't have to be, among the constituents of the state of affairs in question, i.e. among the entities that are talked about and articulated in the content of the statement. The world comes into the picture for purposes of evaluation, but the thoughts that are evaluated need not be metaphysically elaborated thoughts *about* the world. Indeed the users of the language need not even have the ability to entertain such thoughts. Only the theorist needs to be able to talk about the world of evaluation, in her metalanguage.

Modal innocence

Let us consider a simple language without modal operators or other means of talking about worlds ; let us go further and assume that the users of the language don't possess the reflective abilities necessary for thinking about modal issues. They entertain only nonmodal thoughts such as 'Rain is wet'.[20] The possible-worlds semanticist who studies their language will still need to think and talk about the possible worlds relative to which the sentences of the language are evaluated (he may, for example, want to draw a distinction between contingent propositions and necessary propositions) ; but, contrary to what the standard extensional translation suggests, mention of the possible worlds in question will be confined to the theorist's metalanguage.

Now suppose the users of the object-language become sophisticated and start thinking about metaphysical issues. Suppose they come to talk and think about what is *actually* the case as opposed to what *might be* the case. Such modal talk can be formally represented in two ways, as we have seen tensed talk can : by using sentence operators, or by explicitly quantifying world variables in the object-language. If we use the modal framework and introduce modal operators such as 'actually' or 'possibly', *nothing will be changed for the fragment of the language that does not involve those operators.* The sentence 'Rain is wet' will still be a simple,

---

[20] According to some authors there is an implicit modal component in our simplest concepts, so the users of the language we are imagining should be said to lack the reflective abilities necessary for *explicitly* thinking about modal issues. (I am indebted to Robert Brandom for raising this objection.)

modally innocent sentence. The language will simply have been enriched by the introduction of new resources enabling us to construct more complex sentences. But if we use the standard extensional framework and represent modal sentences ('It might be that…', 'Actually…') by means of explicit quantification over possible worlds, as suggested above, then, unless special precaution is taken to avoid that consequence, *a change of language takes place,* not merely an enrichment. In the new language, *all* sentences (including simple sentences) now contain a hidden argument-place for a world. Modal innocence is lost.

I think this move is (almost) as damaging as the previous one — the ascription of thought and talk about possible worlds to modally innocent subjects. Even if the users of the language are sophisticated enough and can think about modal issues, it is misleading to suggest that they *always* think and talk about such issues even when they entertain simple thoughts or utter simple sentences such as 'Rain is wet'. By forcing us to construe e.g. the assertion that rain is wet as involving a covert argument-place which the actual world fills, the extensional translation blurs the cognitively important distinction between the simple, modally innocent assertion 'Rain is wet' and the modal assertion 'Actually, rain is wet'. To maintain that distinction, we have to see modal sentences as *constructed from* simple sentences by the application of modal operators to them. In this way we can analyse the ability to use and understand modal sentences as resting on two distinct abilities : the (modally innocent) ability to use and understand simple sentences ; and the (modally sophisticated) ability to imagine other possible worlds and to contrast the actual world with them. The first ability is independent of the second : we can use and understand simple sentences (e.g. 'Rain is wet') even if we lack the ability to think reflectively about the actual world.

Temporal innocence

The difference we have found between two ways of representing modality can be found also between two ways of representing tense, one which preserves temporal innocence in simple sentences and one which does not.

In tense logic, as we have seen, tense is represented by means of sentence operators. Alternatively, tenses can be represented by adding extra argument-places

(and predicates) for times.[21] If we choose the latter course, it is no longer possible to consider adjectives such as 'warm' or 'yellow' as denoting properties ; they have to be considered as denoting relations — relations between the objects which have the alleged properties and the times at which they have them. As Michael Dummett has pointed out, this relational approach significantly departs from our habitual way of thinking :

> We think of adjectives such as « warm », « smooth », « slender » and so on as denoting *properties* ; properties that a thing may have at one time, and not at another, but nevertheless properties rather than relations between objects and times. And this goes with the way in which we come to understand such adjectives. (…) We do not begin by learning in what relation an object must stand to an arbitrary time for it to be warm or wet at that time, and then, having learned what time is referred to by the adverb « now », derive from this a grasp of what it is for it to be warm now. Rather, we first learn what it is for something to be warm, wet, smooth or slender, that is to say, for the predicate « *is* warm (wet, smooth, slender) » to be applicable to it, where the verb « is » is in the true present tense. From this we advance to an understanding of what is meant by saying of an object that it was or will be warm, etc., at some other time. The advance is made by our acquiring a general grasp of the past and future tenses. That is to say, to understand « was warm » or « will be warm », we apply to our prior understanding of what is meant by saying that something is warm our general comprehension of what it is to speak of how things were or will be at another time. In so doing, we are in effect treating the tenses (and other indications of time) as operators applied to sentences in the present tense of which we have previously acquired an understanding, just as the tense-logical semantics treats them. We could not learn the language in any other way. (Dummett forthcoming : 16-17)

Dummett's complaint about the relational treatment of tenses parallels my complaint about the extensional rendering of modal talk. The relational treatment threatens

---

[21] These are not the only options. Tenses can also be represented as temporal predicates of events in an event-semantical framework.

temporal innocence, just as overt quantification over possible worlds threatens modal innocence.

The present tense

In the temporal case there is a possible objection, due to the fact that tense is (to put it crudely) obligatory in English — or nearly so. Since it is, one may argue that time should not be treated like modality : There are simple, nonmodal sentences, whose characteristics must admittedly be preserved and captured, but there is no such thing as nontemporal talk, hence no such thing as temporal innocence.

From the tense-logical point of view, that objection is misguided. The present tense is not a tense like the past or the future. It is more primitive and, in a sense, temporally neutral. Someone can think 'It is hot in here' even if she has no notion of time whatsoever, hence no mastery of the past and the future. If this is right, mastery of genuine temporal talk rests on two distinct abilities : the ability to use and understand simple, temporally neutral sentences (i.e. sentences in the present) and the ability to think about times and to constrast the past and the future with the present. As in the case of modality, the first ability is independent of the second.

It is true that, when we say or think 'It is hot in here', we talk (or think) about what is presently the case ; we characterize the situation *at the time of utterance*. Yet this is not part of what the sentence itself expresses. The content of the sentence, from the tense-logical point of view, is a function from times to truth-values. When the sentence is uttered, the function is applied to the time of utterance. That is so *whether the sentence is in the present or any other tense*. Even if I say 'It has been hot' or 'It will be hot', I characterize the time of utterance (and, in relation to it, some earlier or later time). The time of utterance, which the sentence is used to characterize, is the time with respect to which we *evaluate* the sentence. The best thing I can do here it to quote Prior :

> If tenses are formed by attaching prefixes like 'It has been the case that' to the present tense, or to a complex with a present tense 'kernel', it is not always true to say that what is in the present tense is understood as a characterisation of the time of utterance ; rather, it characterises whatever time we are taken to by the series of prefixes. The *presentness* of an event, we

may say, is simply the *occurrence* of the event, and that is simply the event itself. But every complete tensed sentence characterises the time of utterance in some way or other, and other times only through their relation to that one. (Prior 1977 : 30)

To sum up, the time of utterance is not represented, it does not feature in the content of tensed sentences ; it only comes into the picture as the circumstance with respect to which the content of a tensed sentence is evaluated. As Prior neatly says in another paragraph, already quoted, on the same page, « tensed propositions are understood as directly or indirectly characterising the *un*mentioned time of utterance ».

Chapter 8
Temporal operators and temporal propositions
in an extensional framework

King's argument rebutted

We have seen that, in the modal/tense-logical framework, modal and tensed sentences are constructed from simple sentences by means of operators. This enables us to construe simple sentences as modally and/or temporally neutral, and thus to preserve modal/temporal innocence. So the dialectic is reversed. Rather than arguing from the presumed existence of temporal operators to the need for temporal propositions, as Kaplan did, we argue from the need for temporal propositions (in order to secure temporal neutrality) to the necessity of construing the tenses as operators. Temporal propositions are required not to make sense of temporal operators (since it is not antecedently obvious that there are such operators), but in order to secure temporal neutrality. And temporal neutrality, which is what temporal propositions give us, is required to capture the fact that our thought and talk is situated in time. The context of our thought/talk fixes the time of evaluation, which is therefore given and does not have to be represented, in contrast to the shifted points of evaluation to which the operators take us.

This does not mean that one cannot 'go extensional'. I said above that, by explicitly quantifying over worlds or times in the object-language, one runs the risk of losing modal/temporal innocence, *unless special precaution is taken*. But, as the qualification suggests, we *can* preserve modal (temporal) innocence even if we opt for an extensional approach. The apparatus of variadic functions presented in Recanati 2002 enables us to do that.

Modal innocence in an extensional framework

In order to preserve modal innocence, only two things are required. First, simple sentences must be modally neutral. Second, modal sentences must be viewed as

constructed from simple sentences by means of sentential operators. None of this precludes us from representing modal talk extensionally if we want to.

In Recanati 2002 I analysed 'Everywhere I go it rains' as resulting from the application of a locative variadic operator to the locationally-neutral sentence 'It rains'. That operator does two things. First, it modifies the adicity of the predicate in the sentence it applies to: it adds an extra argument-place for a location, which can be represented by a free variable. Second, it introduces a restricted quantifier which binds that variable. The operator can be paraphrased as 'for every location *l* such that I go to *l*, in *l* it is the case that'. 'Necessarily it rains' can be represented in the same way, by applying to the modally neutral sentence 'It rains' a sentence operator which explicitly quantifies over possible worlds and can be rendered as: 'for every world *w*, in *w* it is the case that'. Since the variable *'w'* is introduced by the variadic operator, we don't have to treat the emergence of modalities as a radical change in the language, but simply as an enrichment of it; an enrichment which does not affect the simple (nonmodal) sentences, hence preserves modal innocence.

Temporal innocence in an extensional framework

The possibility of such a mixed system shows that there are two orthogonal issues when it comes to time and tense. One has to do with the proper representation of tense : are the tenses best modeled as tense-logical operators, or do they involve the apparatus of (explicit) reference and quantification ? The second issue is the loss of modal and temporal innocence which goes with the adoption of the referential/quantificational framework unless special precautions are taken. The availability of a mixed extensional system with variadic operators shows that, even if the arguments against the modal approach to time and tense are correct, still we can have temporal propositions and operators operating on them. So the standard arguments against the tense-logical approach do not threaten Temporal Relativism.

If we want to represent tense in such a mixed system, we will have to do two things. First, we will have to analyse e.g. the past tense as a variadic operator, 'for some past time *t*, at *t* it is the case that'.[22] (Likewise for the future tense.) Second,

---

[22] 'Past' here means something like 'anterior to the current time', where the current time typically is (but need not be) the time of utterance. If the past tense sentence is in the scope of

sentences in the present will be construed as temporally neutral. This means that we will give the present tense no temporal value whatsoever, in contrast to other tenses such as the past or the future.

This move may sound revolutionary, but in fact the view that the present is temporally neutral is old hat in linguistics. Many theorists, from different traditions, have come up with that idea (see e.g. Malmberg 1979 : 86 for an example from the glossematic tradition). I myself have argued along such lines (Recanati 1995a), and recently, Uli Sauerland has defended that position as well. In a note in *Snippets* 6 (2002), he argues that « the English present tense is semantically vacuous and its interpretive effect is characterized by pragmatic competition with other English tense morphemes, notably the past tense ». This is very Priorean in spirit, even though Sauerland presumably intends to work within an extensional framework with explicit reference to and quantification over times.[23]

Of course, it is not enough merely to say that natural language tenses could be captured within such a mixed system. The work has to be done, and it is only by doing it that one can see the advantages and disadvantages of such an approach. However, my aim in this chapter was very limited. I only wanted to respond to King (2003), by showing that we have a use for temporal propositions, whether one opts for an intensional or an extensional framework.

---

some other temporal operator, then the current time will be the time to which that other operator takes us.

[23] Sauerland's note is a critical amendment to Abusch's theory (Abusch 1997), and Abusch is one of the leading tense theorists who have opted for the extensional approach.

Part Three

The debate over temporalism (2) :
Can we believe temporal propositions ?

Part Three

Chapter 9

An epistemic argument against temporalism

Temporal propositions and belief reports

In 'Temporalism and Eternalism', Mark Richard put forward what many take to be a knock-down argument against temporalism (the view that there are temporal propositions). According to this argument, « the temporalist is unable to give an adequate treatment of attributions of belief » (Richard 1981 : 3).

Richard asks us to consider the following piece of reasoning :

[1] Mary believed that Nixon was president

[2] Mary still believes everything she once believed

Ergo

[3] Mary believes that Nixon is president

As Richard points out, « this argument is not a valid argument in English » and « we ought to reject any position which is committed to [its] validity » (Richard 1981 : 4). Temporalism, Richard claims, is one such position. For the temporalist holds that 'Nixon is president' expresses a temporal proposition p_1, true at any time t iff Nixon is president at t. Let us assume, plausibly enough, that a belief report 'x believes that S' states that the individual referred to by the subject term is belief-related to the proposition expressed by the embedded sentence. It follows that 'Mary believes that Nixon is president' expresses the proposition that Mary believes p_1. On the equally plausible assumption that the past tense in the embedded clause of 'Mary believed that Nixon was president' is semantically vacuous, it follows that [1], 'Mary believed that Nixon was president', expresses the proposition that Mary believed p_1. Now this, together with [2](the proposition that Mary still believes everything she once believed), entails that she still believes p_1, i.e., that she still believes that Nixon is president ! Since that conclusion does not actually follow, there is something wrong with Temporalism.

Richard takes his argument to show that « the objects of belief expressed by sentences are all eternal » (Richard 1981 : 10), i.e., they are classical (Fregean) propositions, not temporal (Aristotelian) propositions. Temporalism can be rescued, Richard points out, if we give up the assumption that « a sentence expresses at most one thing (a proposition) at a time » (1981 : 9). Moderate Relativism as I have described precisely rejects that claim, since it posits two levels of content for every utterance. I will return to Moderate Relativism in the next section. Richard himself describes a view which he calls 'Moderate Temporalism', which rejects the 'single content' assumption :

> We distinguish two different relations of expressing (say, *expresses$_1$* and *expresses$_2$*) and two distinct classes of objects, which we may call contents and propositions. Expression$_1$ is a relation between sentences and contents ; expression$_2$ is a relation between sentences and propositions. Contents may be either eternal or temporal ; propositions are all eternal.
>
> We now take contents to be the bearers of truth and falsity expressed by sentences, propositions to be the objects of belief so expressed. A sentence *S* is true, relative to time *t*, iff there is a content *c* such that *S* expresses$_1$ *c* at *t* and *c* is true at *t*. A sentence *S* expresses, relative to *t*, a belief of a person *u* iff there is a proposition *p* such that *S* expresses$_2$ *p* at *t* and *u* believes *p* at *t*. (Richard 1981 : 10)

In terms of temporally neutral content, we can make sense of the claim that, in a certain sense, two persons who say that it is raining (at different times) 'say the same thing' : that it is raining. And their respective utterances are true iff and only if that constant *lekton* is true at the times of their respective utterances. But what Richard's argument about belief report is supposed to establish is that the content of *belief* is not such a temporally neutral *lekton* : the content of belief is a classical (eternal) proposition. For a moderate temporalist of the sort Richard describes, what is said is a temporal proposition, but what is believed, or what the utterance presents the speaker as believing, is a classical proposition.

49

Richard does not find Moderate Temporalism particularly attractive, because there are utterances like 'What you say is true and I believe it, too' which show that the object of assertion is, or at least can be, the same as the object of belief. At this point, Richard argues, the moderate temporalist will have to distinguish the object of assertion thus understood (a classical proposition, like the object of belief) from 'what the speaker says' in the temporally neutral sense (the *lekton*). But that temporalist notion of 'what the speaker says', distinct both from what the speaker asserts and from what she believes, becomes suspicious, and it is unclear that we need it. « Until some clarification of this notion of 'what is said' by an utterance is given », Richard concludes, « we should remain sceptical » (Richard 1981 : 12).

What a moderate relativist would say

I have spelled out Richard's argument in some detail, in order to make clear where the moderate relativist differs from the hypothetical moderate temporalist described by Richard. Like Richard's moderate temporalist, the moderate relativist distinguishes two types of content and two relations of expressing ; but he would deny that one type of content is what is said, and the other what is believed. The distinction between the two types of content cuts across the distinction between saying and believing. That means that, whether we consider the speaker's assertion or the speaker's belief, we can distinguish two distinct things : the *lekton* (content in the narrow sense) and the complete content or Austinian proposition. Richard himself comes close to that conclusion when, on behalf of the temporal relativist, he draws a tentative distinction between what is asserted (a classical proposition) and what is 'said' (a temporal proposition). In a later paper, he gives examples like

(1)   When Susan saw Kate two winters ago, she swore that Kate was
         pregnant, and when Mindy saw her this Spring, that's what she said too.

and he comments as follows :

> (1) seems to report Susan and Mindy as literally saying the same thing ; if they
> do, presumably they each say something temporally neuter. But… suppose
> that last spring Susan saw Kate and said to herself, '(I guess that) she wasn't

pregnant two winters ago, but she is now'. Then we can surely go to Kate and say,

(2)   When Susan saw you two winters ago, she said that you were
        pregnant, but now she takes that back/denies that/denies what she
        said.

All this, it might be said, suggests what when someone utters a tensed, but temporally unspecific, sentence, two distinct reports of what she said will be possible : one reporting her as having said something temporally specific, and one reporting her as having said something temporally unspecific. And this suggests that utterances of temporally unspecific sentences express, or at least typically express, two things, one temporally unspecific, the other specific. (Richard 2003a : 39-40)

That is exactly what a moderate relativist will say ; and the moderate relativist will point out, as Richard himself does in the later paper, that the same point can be made with respect to belief.

Richard gives the following example of a *belief* ascription where the object of belief seems to be a temporal proposition :

(3)   Bob went to the monkey house, and now he thinks that he's been infected
        with the Ebola virus. Every time he goes there he thinks that ; he's
        convinced one of the monkeys is a carrier.

The word 'that' in 'Everytime he goes there he thinks that' seems to refer to a temporal proposition (since the eternal proposition believed by Bob after visiting the monkey house changes from one visit to the next). Richard, however, thinks the evidence is misleading. He has a story to tell regarding examples like (3), which story does not appeal to temporal propositions as objects of belief.[24] His reason for

_____

[24] According to Richard (2003a : 41-2), in 'Every time he goes there he thinks that', 'that' does not refer to a temporal proposition denoted by the antecedent 'that'-clause 'that he has been infected with the Ebola virus'. Rather than construe 'that' as a device of cross-reference,

resisting the view that there are two possible objects of belief, corresponding to the two levels of content distinguished by the moderate relativist, is that « diachronic agreement or disagreement seems to be, of necessity, a matter of agreement or disagreement about something temporally specific » (Richard 2003a : 40). So we are back to Richard's original argument : whether one changes one's mind or retains one's belief is a matter of still believing (or ceasing to believe) the same classical propositions. *When it comes to assessing inter- or intra-individual (dis-)agreement, only classical contents count.* So if we know that Mary retained all of her previous beliefs, we will not conclude that she still believes that Nixon is president even though we know that, twenty years ago, she believed that Nixon was president. On this issue, Richard has not changed his mind from 1981 to 2003, and his objection to temporal propositions as the objects of belief still stands.

---

we may, « with a fair amount of plausibility », construe it as a device of ellipsis, Richard says (p. 42). On that analysis 'Every time he goes there he thinks that' is *short for* 'Every time he goes there he thinks that he has been infected with the Ebola virus', and in *that* sentence the belief that is ascribed to Bob is temporally specific (eternal) rather than temporally neutral. The sentence says that for every t such that Bob visits the monkey house at t, he comes to believe (after t) that he has been infected with the Ebola virus *at t*.

Chapter 10
Rebutting Richard's argument

The duality of belief
(and the ambiguity of belief reports)

In response to Richard, the moderate relativist can point out that belief reports have (at least) two distinct functions. First, one may report someone's beliefs in order to assess their (dis)agreement with the facts, or with the beliefs of other people, or the beliefs held by the same person at different times, about the same facts. Such belief reports will typically focus on the truth-conditional properties of the belief, hence on its complete content. Second, one may report someone's beliefs in order to draw some 'internal' connection between those beliefs and other states or acts of the same person, for example her sensory experiences, her actions, or other beliefs potentially or necessarily held by her. Such belief reports focus not on the referential properties of the belief but on what McGinn calls its 'intra-individual causal-explanatory role'. This distinction between two functions of belief reports is well-known and it has been extensively documented in the seventies. In his classical paper on these issues, McGinn writes :

> Our concept of belief combines two separate elements, serving separate concerns : we view beliefs as causally explanatory states of the head whose semantic properties are, from that point of view, as may be ; and we view beliefs as relations to propositions that can be assigned referential truth-conditions, and so point outward to the world. This bifurcation of content can be seen as stemming from the point that beliefs involve internal representations, and these inherently present a dual aspect. (McGinn 1982 : 216)

Since there is this duality in our notion of belief, it is not surprising that there is an ambiguity in a belief report like 'Susan believes that Kate is pregnant'. This may

ascribe to Susan either the internal state of believing Kate pregnant, a state one may be in at different times ('relativist' interpretation); or it may ascribe to her a belief with a certain truth-conditional content, which content depends, as we have seen, upon external factors such as the time at which the belief is held ('classical' interpretation). On the latter interpretation, Susan's belief can change from one occurrence of the internal state to the next, even though the internal state itself does not change. At *t*, Susan is in the state of believing Kate pregnant, and she thereby believes the classical proposition that Kate is pregnant at *t* ; at *t'* Susan is in the same state, but the classical proposition she now believes is the (distinct) proposition that Kate is pregnant at *t'*. If, on the classical interpretation, we say that someone's beliefs have not changed, it follows that she believes all the *classical* propositions she formerly believed ; but it does not follow that her internal doxastic state has not changed. On that interpretation the argument Richard presents as invalid is indeed invalid. From the fact that, at a certain time *t*, Mary was in the state of believing Nixon president, and thereby believed the classical proposition that Nixon is president at *t*, plus the fact that she still believes all the classical propositions she once believed, it does not follow that she still is in the state of believing Nixon president and thereby believes the classical proposition that Nixon is president now.

So, on the classical interpretation of a belief report, the object of belief is indeed the complete content.[25] But that is not the only possible reading of a belief

---

[25] In the SMR framework, the complete content of an utterance or thought is the Austinian proposition, which determines the representation's truth-conditions. But different Austinian propositions can have, or determine, the same truth-conditions, and what matters for (dis)agreement are only the truth-conditions. Two people with two different perspectives on some fact may nevertheless agree or disagree : the complete contents of their beliefs are distinct Austinian propositions, but that over which they (dis)agree is more abstract and corresponds to an *equivalence class* of Austinian propositions. That is what I have in mind when I talk of the content of belief as being a 'classical proposition'. (Classical propositions abstract from the perspective which Austinian propositions encode.) In other words, I reject the assumption that doxastic agreement requires *identity* of belief content. We can define (dis)agreement as follows : Two persons (or the same person at different times) agree iff the Austinian propositions which is the complete content of their beliefs belong to the same

report. There is another reading, where the ascriber is interested in the intra-individual causal-explanatory role of the ascribed belief. In such cases what matters is the *lekton*, not the complete truth-conditional content.

The lekton as cognitive content

That we need the *lekton* in such cases has been forcefully argued by Barwise, who gives the following example. Suppose Holmes and Watson face each other. In between stand the salt and the pepper. Holmes says 'The salt is left of the pepper', because the salt is left of the pepper from Holmes's perspective. From Watson's perspective, the pepper is left of the salt; however, Watson is mistaken as to which shaker is which, and he wrongly says 'The salt is left of the pepper'. Holmes and Watson apparently 'say the same thing' (so they express the same *lekton*) but Holmes is right and Watson wrong (so they believe different classical propositions, or different Austinian propositions, because they each relativize the *lekton* to their own perspective). In the classical framework, Barwise points out,

> we have nothing in the theory that classifies the similarity in attitudes of Holmes and Watson in cases like these. And it is this similarity that leads them to make the same bodily movements, reaching in the same direction, though toward different objects, when they want the salt. (Barwise 1989a: 240).

In other words, if what we are interested in is the state Holmes and Watson are both in, and the causal-explanatory role of that state, then we should accept that (in the relevant sense) they believe the same thing : they both have a belief with a certain *lekton* as content, which *lekton* determines different truth-conditions when evaluated with respect to their distinct perspectives. So there is a sense in which Holmes and Watson believe the same thing in that situation, and there is also a sense in which they do not believe the same thing. When arguing that the content of belief must be eternal (classical), Richard simply focusses upon the sense which is relevant to belief

equivalence class from the truth-conditional point of view. In such a case they can be said to believe the same classical proposition.

reports whose function is to assess (dis)agreement with the facts ; but everybody knows that that is not the sole function of belief reports.

Back to Richard's argument

If I am right, shouldn't there be an interpretation in which the argument Richard discusses *is* valid ? Let us reconsider that argument :

[1] Mary believed that Nixon was president
[2] Mary still believes everything she once believed
Ergo
[3] Mary believes that Nixon is president

Since [3] does not follow, Richard argues that [1] cannot be understood as describing a certain relation between the believer and a temporally indefinite *lekton* ; for if it did, [3] would follow. Now I hold that [1] *can* indeed describe a relation between Mary and the temporally indefinite *lekton*. That is the 'relativist' interpretation of the belief report. On that interpretation, [3] ought to follow — but it does not. Does this not show that Richard is right, and the moderate relativist wrong ?

I do not think so. The reason why [3] is hard to accept in this context may be due to the fact that the universal quantification over beliefs in premiss [2] has to be understood either as quantifying over complete contents or as quantifying over *lekta*, since those are different sorts of things. Let us assume that a choice has indeed to be made from the outset when interpreting [2]. Then, arguably, the 'classical proposition' interpretation is more salient, for some reason, and that is why [3] does not follow from [1] and [2] even though [1] *can* be interpreted as describing a relation between the subject and a temporally indefinite *lekton*. To check that this is correct, we have only to rephrase [2] so as to make the *lekton* interpretation more salient.

To show this, let us go back to the Susan/Kate example and run the Richard argument in that context, while suitably modifying (the counterpart of) premiss [2]. We get :

[1'] Susan believed that Kate was pregnant

[2'] Susan is exactly in the same doxastic state she was in. (Imagine a Rip van Winkle context in which, unbeknown to her, she has just awoken from a two-year sleep.)

[3'] *Ergo* : she still believes that Kate is pregnant.

This argument seems valid. Now the only substantial thing I have done (besides transposing the example) is change the second premiss in order to force the causal-explanatory reading of the first premiss. That is precisely the reading which Richard wrongly treats as nonexistent. So my example shows that, *pace* Richard, *there is* a reading in which the embedded clause in a belief report like [1] or [1'] is used to pick out the *lekton* rather than the complete content of the ascribed belief.

Chapter 11
Relativistic disagreement

Shared circumstances

So far, I have conceded that when the point of a belief report is to assess (dis)agreement with the facts or with other people or with oneself at different times, it is the complete content of the belief that matters, since the complete content is what determines possible-worlds truth-conditions. But this too can be disputed.

There are two types of example which cast doubt on the idea that agreement or disagreement can only be about complete contents. The first type of example involves cases in which all the parties to the conversation are in the same situation — for example, they share their location. In such cases, there is no objection to one of them asserting a place-neutral proposition such as 'Sidney is nearby', true at any place *l* iff Sidney is near *l*. (I borrow the example, and the argument, from Egan forthcoming.)

Let us assume the belief-transfer model of assertion, according to which the function of assertion is to transfer belief in the asserted proposition (Stalnaker 1978). In the present case the transfer can be described as follows :

> Speaker A, at place *l*, accepts the place-neutral proposition that Sidney is nearby, and thereby believes the classical proposition that Sidney is near *l*.

> Speaker A asserts the place-neutral proposition that Sidney is nearby.

> As a result, audience B, also at place *l*, comes to accept the place-neutral proposition that Sidney is nearby, and thereby believes the classical proposition that Sidney is near *l*.

Even though what is asserted, and what is transferred, is a place-neutral proposition, the audience comes to believe the same classical propositions as the speaker. Since

their situations are the same, the *lekta* cannot determine different truth-conditions with respect to their respective situations. That is why the belief transfer can take place directly at the *lekton* level, in this type of case. So it is not true that the content of assertion can only be the complete content. In shared-situation cases, the *lekta* go proxy for the complete contents (Barwise 1989a : 253), and we can assert them (Egan forthcoming).

Just as we can assert *lekta* in shared-situation cases, we can agree or disagree about them. The function of assertion is to transfer belief : the speaker says something, and the audience is supposed to accept what the speaker has said. Sometimes, however, the speaker says something and the audience does *not* accept what the speaker has said : they disagree. Like assertion, disagreement, in shared-situation cases, can be over *lekta*. I say 'Sidney is nearby', and you respond 'No, it isn't'. If our locations were different, there would be no point in so disagreeing about *lekta* : our respectively accepting and denying the place-relative proposition that Sidney is nearby would not entail any truth-conditional incompatibility between our beliefs. But genuine (dis)agreement over *lekta* is possible whenever the situation of evaluation is shared.

This straightforwardly applies to the temporal case : there can be genuine agreement or disagreement over temporally neutral propositions whenever the time of evaluation is shared. If it is not — as in cases of diachronic disagreement — then the content in dispute must be temporally specific. Here, Richard seems to be right : « diachronic agreement or disagreement seems to be, of necessity, a matter of agreement or disagreement about something temporally specific » (Richard 2003a : 40).
Even that has been disputed, however. This brings us to the second type of example which casts doubt on the view that, in matters of agreement or disagreement, only complete contents count.

Faultless disagreement

Sometimes, we genuinely disagree about a certain *lekton*, even though (it seems) we are not in the same situation. Thus, looking at a painting, I say : 'This is beautiful'. You disagree : 'No, it's ugly'. In a sense, we are both right, since for me it is beautiful, while for you it is ugly ; but we disagree nonetheless. Or consider epistemic modals. I

say 'The treasure might be under the palm tree'. I am right since, for all I know, the treasure might be there — nothing in my epistemic state rules out the treasure's being there. Later, however, I learn that the treasure is not on the island (where the palm tree is). This rules out the treasure's being under the palm tree, and in my new epistemic situation, I assert : 'The treasure cannot be under the palm tree'. Again, I am right since, in my new epistemic situation, there is something that rules out the treasure's being under the palm tree. What is strange, however, is that I can now disagree with my former self. I can say : 'I was wrong — the treasure cannot be under the palm tree'. How can that be ? If I was right, given my epistemic situation then, how can I later judge that I was wrong ?

Such cases of faultless disagreement suggest that sometimes at least, agreement or disagreement is about the *lekton*, even though the disagreeing parties evaluate the *lekton* with respect to distinct situations/perspectives. This makes sense if one is a radical relativist. For a radical relativist, the *lekton* is complete. It *is* the content – that which one asserts, believes, and over which one agrees or disagrees with others.[26] The situation of evaluation is not an aspect of content (broadly understood) but something external to content. As Prior puts it,

> Aristotle… says that 'statements and opinions' vary in their truth and falsehood with the times at which they are made or held, just as concrete things have different qualities at different times ; though the cases are different, because *the changes of truth-value of statements and opinions are not properly*

---

[26] MacFarlane protests that in many cases, *he* does not take the *lekton* to be complete, even though he is a radical relativist : « So, for example, I could say (with the temporalist) that the time of utterance is not part of the *lekton*, but rather part of what a use of the *lekton* concerns, and still tell my story about the assessment-sensitivity of future contingents» (MacFarlane, p.c.). But MacFarlane's Radical Relativism is not absolute, as we have seen (footnote 6) : MacFarlane is a radical relativist with respect to some examples but not with respect to others. My claim is conditional : *if* one is a radical relativist with respect to a given type of sentence, e.g. 'The treasure may be under the palm tree', whose truth-value is relative to something in addition to a possible world (in this case, an epistemic state), *then* one holds that the *lekton*, that is, the content of the sentence independent of that thing, is complete and can be the object or assertion, belief, or (dis)agreement.

speaking changes in these statements and opinions themselves, but reflections of changes in the objects to which they refer (a statement being true when what it says is so, and ceasing to be true when that ceases to be so. (Prior 1967 : 16 ; emphasis mine)

So a radical relativist has a story to tell about faultless disagreement, and that involves giving up the claim that agreement or disagreement in non-shared situations is of necessity a matter of agreement or disagreement about classical content. (Ironically, Mark Richard is among the recent advocates of Radical Relativism.)

Faultless disagreement in a Moderate relativist framework :
The Kölbel-Lasersohn view

Some authors (Kölbel forthcoming ; Lasersohn 2005) have claimed, or implied, that faultless disagreement merely calls for a distinction between the content over which we agree or disagree (the *lekton*) and the utterance's possible-world truth-conditions. The latter depend upon, and covary with, the situation of evaluation. On this view the moderate relativist framework, with its distinction between two levels of content, is sufficient to account for faultless disagreement. Indeed faultless disagreement is the best evidence we have for Moderate Relativism : the phenomenon demonstrates that we need one level of content as bearer of (unrelativized) truth-value, and another level of content as the object over which rational subjects (dis)agree.

I think these authors are mistaken. There are plenty of cases in which we must distinguish between the *lekton* and the complete content, but in which there can be no genuine (dis)agreement about the *lekton*. Thus I call you on the phone, and commenting upon my situation I say 'It is raining'. If you say 'No, it isn't', meaning that there is no rain in your situation, there is misunderstanding rather than genuine disagreement. Or, adapting Barwise's example, suppose that Watson says 'The salt is left of the pepper', and Holmes, speaking from his own perspective, replies 'No it is not'. Clearly, there is no substantive disagreement here. If each of them is talking about his own perspective, there is misunderstanding rather than genuine disagreement. The same considerations apply to the temporal case. At time *t*, you say 'It is raining'. Later, when the sun is shining again, you say 'It is not raining'. You

cannot conclude 'so I was wrong'. Here, as Richard points out, genuine disagreement can only be about temporally specific contents.

So Moderate Relativism by itself, with its two levels of content, does not give a solution to the problem of faultless disagreement, contrary to what Kölbel and Lasersohn believe. This does not mean that no solution is available within the moderate relativist framework, however. I think a solution is available, along the following lines.[27]

Faultless disagreement in a Moderate relativist framework :
A better view

The statement that the painting is beautiful must be evaluated with respect to some esthetic standard, but that standard cannot be the speaker's esthetic standard ; for if that were the case, then the speaker's 'It is beautiful' and the audience's reply 'No, it is ugly' would not contradict each other. When the speaker says 'It is beautiful', he means something stronger than merely 'It is beautiful by my standards' or 'I find it beautiful'. To be sure, when realizing that the audience does not share his taste, the speaker may retreat to the weaker statement : 'I find it beautiful', and this gives us a hint as to what the proper analysis of the stronger, unqualified statement should be. I suggest the following : 'It is beautiful' means that it is beautiful 'for us', that is, for the community to which the speaker and his audience belong. When the audience says 'No it isn't', the speaker realizes that the stronger statement is incorrect (since the audience does not actually find the painting beautiful), and he retreats to the weaker statement.

Lasersohn has an objection to that analysis. Were it correct, he says, it would not be possible for the speaker to maintain the stronger statement and keep disagreeing with his audience. But this is clearly possible, as in the following exchange:

— This painting is beautiful.

---

[27] I will discuss only a case involving standards of taste. Whether the type of solution I put forward extends to other cases (such as epistemic modals) is an issue I will leave for further research. (See von Fintel and Gillies 2004 and 2006 for a step in that direction.)

— No, it's ugly.

— I tell you it is beautiful.

To handle this type of counter-example, we must introduce a certain flexibility as to what counts as 'the community' from whose point of view the *lekton* is meant to be evaluated. I may judge that my audience deviates, by her bad taste, from the esthetic standards of the community to which we both belong, those standards being fixed by e.g. the community's experts. This enables me to disagree with my audience, even though what I claim is that the painting is beautiful *for us*, i.e., for our community.

Whether, after the hearer has expressed disagreement, the speaker retreats to the weaker statement ('I find it beautiful') or not ('I tell you it's beautiful !'), the right conclusion to draw is that there is *no* genuine faultless disagreement in those cases. The speaker who retreats to the weaker statement realizes that he was *mistaken* when he assumed that the painting was beautiful for both him and his audience ; so he was at fault. And when the speaker refuses to retreat and maintains his position, because he « judges that his audience deviates, by her bad taste, from the esthetic standards of the community to which they both belong, those standards being fixed by e.g. the community's experts », he finds *her* at fault. In both cases, there is genuine disagreement but it is not faultless (Stojanovic forthcoming). Nor are those cases cases in which the speaker and his audience evaluate the *lekton* with respect to different standards: they both appeal to the standards of the community to which they belong.

Those, however, are the simple cases. More difficult to handle is the following case (inspired from an example by Lasersohn). Suppose the speaker finds the painting beautiful, *and no one else (and in particular, no expert) does*. Suppose, moreover, that the speaker knows that he is alone in finding the painting beautiful. Still, as the opinionated person he is, he maintains : 'This is beautiful !' Clearly, he does not mean simply that *he* finds the painting beautiful. He refuses to retreat, and maintains the stronger claim. Can we handle that example consistently with the view that 'This is beautiful' means *This is beautiful for us* ? Can we maintain that the speaker, even in such a case, appeals to the standards of the community ?

I think we can. As Johan Brännmark writes (in connection with moral judgments),

When speaking for a collective, I cannot deviate from its present view without ceasing to speak for the collective. But communities are multi-generational and by their very nature they always have one foot in the future... So I think we can distinguish between two ways in which we can speak for the community : first, there is the purely representative one ; second, there is a progressive one. In passing progressive moral judgment *I am deviating from the present community in the direction towards which I find that the community would move if the people in it thought things through really well.* (Brännmark forthcoming : 18 ; emphasis mine).

In this sort of case — when no appeal to experts is relevant and each of the two disputants is expressing his or her own taste and trying to impose it — it makes sense to say that they are 'both right', even though they disagree : for they are both entitled to interpret the community's standards the way they want in passing progressive judgment. This freedom explains how they can disagree about the truth-value of 'This is beautiful' even though, in evaluating the statement, *they both appeal to the standards of the community*. This is possible because the standards in question are not fixed once for all — they are up to us. Each of us can contribute to shaping them, and this is what we do in this type of case when we say 'It is beautiful' or 'It is ugly'. Austin would perhaps make this point by saying that there is a performative element in such statements.

If we accept this explanation of the phenomenon, we can maintain that agreement or disagreement is about complete contents,[28] except when the situation of evaluation is shared (in which case the *lekton* can go proxy for the complete content). Alleged 'faultless disagreement' cases such as those we have discussed are no exception. In such cases, the disagreement is about the Austinian proposition consisting of the *lekton* together with the standards of the community.[29] The speaker says that the painting is beautiful (for the community), and the audience denies that it is beautiful (for the community). The disagreement ultimately bears upon what the community standards are, or should be. If the disagreement bears upon what the

---

[28] In the sense which has been defined in footnote 25.

[29] Since both disputants implicitly appeal to the standards of the community, we can also treat that case as a case in which the situation of evaluation is shared.

community standards *are*, it is not faultless. If, as in the last case we considered, it bears upon what the standards *should be*, then it is, arguably, faultless, but even in that case the disagreement is over the complete content.

Part Four

Relativization and Indexicality

Chapter 12

Index, context, and content

Indexicality : broad and narrow

There is one form of relativization we have not talked about so far. A sentence such as 'I am French' is true with respect to certain contexts (contexts in which it is uttered by a French person) and false with respect to other contexts. A sentence that has this property is said to be *indexical.* This is similar to *contingency* (the fact, for a sentence, of being true with respect to some possible worlds and false with respect to others) and to *temporality* (the fact of being true with respect to certain times and false with respect to others). But how similar ?

In the article which bears the same title as this chapter, Lewis claims that indexicality is a general phenomenon, of which the other forms of relativity are particular species. He writes :

> It is a feature of any context, actual or otherwise, that its world is one where matters of contingent fact are a certain way. Just as [because of indexicality] truth-in-English may depend on the time of the context, or the speaker, or the standards of precision, or the salience relations, so likewise may it depend on the world of the context. Contingency is a kind of indexicality. (Lewis 1980 : 82)

The idea is that, when the truth-value of an utterance depends upon a circumstantial feature such as a world, a time, a place, or whatever, what fixes the relevant world, time, place or whatever and makes evaluation possible is the context of utterance. The world, time, place etc. is the world, time, place, etc., of the context, that is, the world, time, place, etc. at which the sentence is uttered.

But is this true ? There are two types of counterexample we should consider. First, an utterance like 'It is raining' may well concern a place distinct from the place of utterance. If the speaker and the hearer are involved in a conversation about what is going on at a certain place *l*, distinct from the place of utterance, the sentence 'It is

raining' uttered in that context will presumably be about *I*, not about the place of utterance. To be sure, in such a case, it is still the context that determines the place which is talked about, viz. *I*. But the place in question is not *the place of the context*.[30] Here we see a difference between the place-relative sentence 'It is raining' and 'It is raining here'. While the first sentence expresses a place-relative proposition, true at an arbitrary place iff it is raining at that place, 'It is raining here' can only be evaluated with respect to the place of utterance. This casts doubt on the idea that e.g. place-relativity is a form of indexicality. It would seem, rather, that indexicality is a (very) special form of relativity.

Be that as it may, I will temporary ignore this difference between circumstance-relativity and indexicality. Going along with Lewis, I will pretend that an utterance of 'It is raining', just like an utterance of 'It's raining here', is true iff it is raining at the place of utterance.[31] On this view, circumstance-relativity *is* a form of indexicality.

Still, 'It is raining' and 'It is raining here' behave differently when embedded, and this provides the second type of counterexample to the claim that circumstance-relative sentences are evaluated with respect to the circumstance of the context. Even if we hold that an utterance of the sentence 'It is hot' is true iff it is hot at the place (and time) of utterance, we have to acknowledge that this is no longer true if we embed the sentence. When such a sentence falls in the scope of a temporal or locational operator it is evaluated with respect to the time or place to which the operator takes us, not the time or place of utterance. Similarly, a sentence may be evaluated with respect to an arbitrary world, not necessarily the world of the context. That is what happens when the sentence falls within the scope of a modal operator : the sentence is then evaluated with respect to the world(s) to which the operator takes us.

The possibility of shifting the point of evaluation through the use of operators forces us to introduce a distinction between sentences like 'It is hot', which are

---

[30] Similarly, in an appropriate context, a simple sentence such as 'It is raining' may well concern a hypothetical world rather than the actual world. The world an utterance concerns is determined by context but it need not be the world of the context.

[31] Lewis's claim will be scrutinized (and rejected) in Chapter 38. See also Chapter 36 on the difference between situation-relativity and context-relativity.

indexical in the broad sense talked about so far, and sentences like 'It is hot here and now', which are indexical in a narrow sense. The former can be evaluated with respect to an arbitrary time and place if we embed it under an operator such as 'somewhere it will be the case that' ; not so with the latter. Indexical adverbs like 'here' and 'now' keep the evaluation relative to the context of utterance, *even when the sentence in which they occur is embedded*. Similarly, an indexical sentence such as 'Actually S' can only be evaluated with respect to the world of the context because the indexical adverb 'actually', like the indexical adverb 'now', keeps the evaluation relative to the context of utterance.

Lewis was aware of all this. He took the dependence of truth-value on the context of utterance (indexicality in the broad sense) to be the basic phenomenon, and introduced *another* phenomenon — shiftiness — which accounts for the difference between sentences that are indexical in a strict sense and sentences that are not.

Shifting and freezing

Some features of context are 'shiftable', Lewis says. Modal and temporal operators shift the point of evaluation : they take us to a point of evaluation distinct from the initial or default point which is fixed by the context of utterance. Thus a sentence in the scope of a temporal operator is evaluated with respect to a time distinct from the time of utterance. When the point of evaluation is distinct from the context of utterance because some feature of the context has been shifted by an operator, the sentence in the scope of the operator is evaluated with respect to an 'index' (a *n*-tuple of features) that is not the 'index of the context', Lewis says.

Not all features of the context can be shifted. There are temporal operators, locational operators, modal operators, and so on and so forth, which shift the time, location, and world feature of the context respectively; but there are many features of the context to which some expression may be sensitive but which no operator can shift. When a feature is not shiftable, a sentence which contains an expression sensitive to that feature of context will be indexical in the narrow sense : its evaluation will depend upon the context of utterance, whether it is embedded or not.

In some cases a feature of context is shiftable (because there is an operator which can shift it), but it is possible to 'freeze' it so as to keep it unshifted, or rather,

so as to 'un-shift' it. This is what indexical expressions — expressions such as 'here' or 'now' — typically do. With respect to the relevant feature, they constrain the evaluation index to be the index of the context. Even if a sentence is embedded under a temporal operator which shifts the time of evaluation (as e.g. 'it will soon be the case that'), the presence of a temporal indexical in that sentence takes us back to the time of utterance. Kaplan gives the following example :

Suppose I say, at t_0, « It will soon be the case that all that is now beautiful is faded. » Consider what was said in the subsentence,

All that is now beautiful is faded.

I wish to evaluate that content at some near future time t_1. What is the relevant time associated with the indexical 'now' ? Is it the future time t_1 ? No, it is t_0, of course : the time of the context of use. (Kaplan 1989 : 498)

Such is the freezing power of indexicals : in 'It will soon be the case that all that is now beautiful is faded', 'now beautiful' means *beautiful at the time of utterance,* even though the embedded sentence in which that expression occurs is evaluated with respect to a future time (the time to which 'It will soon be the case that' takes us). As Prior writes,

The essential point about… 'now' is that however oblique the context in which it occurs, the time it indicates is the time of utterance of the whole sentence. (…) In 'it will be the case that I am sitting now', the word 'now' indicates the same time that it would indicate if it occurred in the principal clause — the time of utterance. (Prior 1968b/2003 : 174)

To sum up, if we follow Lewis and take indexicality (in the broad sense) to be the most general phenomenon, we will have to distinguish within the class of broadly indexical expressions — expressions whose evaluation is relative to some feature of context — two sub-classes : on the one hand, the expressions whose evaluation is relative to a feature of context that can be shifted ; and on the other hand, the expressions whose evaluation is relative to a feature of the context which, for one

reason or another, cannot be shifted. (The latter are indexical in the narrow sense.) Unshiftability itself will come in two varieties : some features of context are unshiftable because no operator exists in the language which has the power of shifting it, and others resist shifting because, although there are operators which can shift them, there are also freezing devices which prevent this from happening (or rather, which take us back to the index of the context, when that happens). Indexical adverbs like 'here' and 'now' have this freezing power, in contrast to merely time-relative expressions (e.g. the tenses) or merely location-relative expressions (e.g. 'to the left').

Two-dimensional semantics

Lewis concludes that we must distinguish

> two different dependencies of truth on features of context : context-dependence and index-dependence. We need the relation : sentence *s* is true at context *c* at index *i*. We need both the case in which *i* is the index of the context *c* and the case in which *i* has been shifted away, in one or more coordinates, from the index of the context. (Lewis 1980 : 88)

In other words, because of the double dependence of truth-value on context and index, we need a two-dimensional semantics. As Lewis emphasized, there are two ways to achieve that goal.

The first approach consists in assigning to sentences, as their meanings, rather complicated functions from context-index pairs to truth-values. A sentence like 'It is raining here' will be true with respect to a context of utterance *c* and an index *i* if and only if it is raining at the place of *c* at the time of *i* in the world of *i*. In contrast, 'It is raining' will be true with respect to a context *c* and an index *i* if and only if it is raining at the place, time and world of *i*, and 'It is actually raining' will be true at a context *c* and an index *i* if and only if it is raining at the place and time of *i* in the world of *c*.

In this framework, the evaluation process for a given utterance will start from the context of utterance and the index of that context. All index-relative elements will be assigned values relative to that index, until an index-shifting operator is met ; the

current index will then shift (under the relevant feature — that which the operator concerns) and index-relative expressions processed thereafter will be evaluated with respect to the shifted index, until another shift occurs... As for indexicals (in the narrow sense), they will always be evaluated with respect to the context of utterance and will therefore be indifferent to the current status of the index.

The other approach consists in assigning to each sentence two simple semantic values instead of a complicated one. The first semantic value is what Kaplan calls the sentence's *content*, which depends upon the context. The content of a sentence results from evaluating the indexicals in it with respect to context. What we end up with, when that is done, is a simple function from indices to truth-values. So the sentence 'It is hot in here', when evaluated with respect to a context c, expresses the temporal but location-specific proposition that it is hot at the place of c. That proposition is the kaplanian 'content' of the utterance, and it is true at an arbitrary index i consisting of a world w and a time t iff it is hot at the place of c in w at t. The sentence 'It is hot' will have a different content : it will express a location-indefinite proposition true at an index iff it is hot at the world, time *and* place of that index.

Kaplanian contents are what Lewis calls 'variable but simple semantic values'. They are contextually variable, in contrast to the semantic values of the first approach (functions from context-index pairs to truth-values) ; and they are simpler since the argument of the function is simpler. But there is no overall gain in simplicity, for Kaplan must assign sentences another semantic value to deal with their constant linguistic meanings. That second semantic value is their *character*, which Kaplan models as functions from contexts to contents (i.e., functions from contexts to functions from indices to truth-values).

Which solution is better ? According to Lewis, neither is better than the other. The two solutions are simply equivalent : whatever one can do, the other can do too. Both incorporate the double dependence of truth-value on index and context, and that is all that matters.

Chapter 13
The two-stage picture :
Lewis vs Kaplan and Stalnaker

According to Lewis, we need a two-dimensional semantics, but there is no reason to opt for a two-stage picture such as Kaplan's. It is sufficient to assign to sentences functions from context-index pairs to truth-values. Now I myself have opted for the two-stage framework in this book. Following Barwise, I have taken the *lekton* to be, not the meaning of the sentence-type, but the relativized proposition that results from assigning contextual values to indexicals ; and I have emphasized how similar the *lekton* thus understood is to Kaplan's 'content'. I must therefore embark upon a detailed examination of Lewis's argument with Kaplan and Stalnaker over the two-stage picture. Can we, or can we not, provide a justification for that framework, and for the notion of *lekton* as it has been construed so far in this book ?

Kaplan

Kaplan claims that his notion of content captures the intuitive notion of 'what is said' by an utterance. If I say 'I am bald', and, expressing your agreement, you say to me, 'Indeed, you are bald', we both *say the same thing* ; and that corresponds to the 'content' in the two-stage picture.

　　　Lewis responds that there is no point in capturing 'the' everyday notion of 'what is said', since it is far from univocal :

> It can mean the propositional content in Stalnaker's sense… It can mean the exact words. I suspect that it can mean almost anything in between. (Lewis 1980 : 97)

We may interpret Kaplan differently, however. Maybe what he meant was simply that his notion of content (functions from indices to truth-values) captures *the content of*

speech acts, and more specifically, the content of locutionary acts : what someone says (or states) when he or she utters a declarative sentence such as 'You are bald' or 'I am bald'. Any theory of language has to show how locutionary content is derived from the conventional meaning of the sentence and the context. While Kaplan's 'characters' correspond to the conventional meanings of sentences, his 'contents' arguably correspond to the contents of the locutionary acts performed by uttering those sentences. This role for contents would provide some kind of justification for the two-stage approach.

<center>*Stalnaker*</center>

However we interpret Kaplan, the type of justification I have just mentioned is quite explicit in Stalnaker's early writings in defense of the two-stage picture. The following quotations are from his paper 'Pragmatics' (Stalnaker 1970 : 277-8) :

> The scheme I am proposing looks roughly like this : The syntactical and semantical rules for a language determine an interpreted sentence or clause ; this, together with some features of the context of use of the sentence or clause determines a proposition ; this in turn, together with a possible world, determines a truth-value . An interpreted sentence, then, corresponds to a function from contexts into propositions, and a proposition is a function from possible worlds into truth-values.
>
> According to this scheme, both contexts and possible words are partial determinants of the truth value of what is expressed by a given sentence. One might merge them together, considering a proposition to be a function from context-possible worlds (call them points of reference) into truth values… [This] is a simpler analysis than the one I am sketching ; I need some argument for the necessity or desirability of the extra step on the road from sentences to truth values. The step is justified only if the middlemen — the propositions — are of some independent interest.

Stalnaker then proceeds to give the following argument for the desirability of the extra step :

The independent interest in propositions comes from the fact that they are the objects of illocutionary acts and propositional attitudes. A proposition is supposed to be the common content of statements, judgments, promises, wishes and wants, questions and answers, things that are possible or probable. The meanings of sentences, or rules determining truth values directly from contexts, cannot plausibly represent these objects.

If O'Leary says « Are you going to the party ? » and you answer, « Yes, I'm going », your answer is appropriate because the proposition you affirm is the one expressed in his question. On the simpler analysis, there is nothing to be the common content of question and answer except a truth value. The propositions are expressed from different points of reference, and according to the simpler analysis, they are different propositions. A truth value, of course, is not enough to be the common content. If O'Leary asks « Are you going to the party ? » it would be inappropriate for you to answer, « Yes, snow is white. »

So, for Stalnaker, propositions are needed as semantic values intermediate between constant meanings and truth-values because they serve as contents of illocutionary acts and propositional attitudes. What are we to think of this argument ?

A dilemma for the two-stage theorist

According to Lewis, Stalnaker's argument is no good, for two reasons :

(1) Within the framework of the single-stage approach, we can easily define the proposition expressed by an utterance with respect to a given context *c.* (It is the proposition that is true at any world *w* iff the sentence is true at *c* and the index of *c* when the world of that index has been replaced by *w*.) So it is not true that if we opt for the single-stage approach, there will be 'nothing to be the common content of question and answer except a truth value'. The relevant notion of content will be extractable, even if it is not as immediately available as it is in the two-stage picture.

(2) In any case, Stalnaker's propositions cannot be assigned to sentences as their semantic values (or one of their semantic values), because they do not satisfy the

principle of compositionality. That they are not compositional is easy to see on an example :

> The propositional content of 'Somewhere the sun is shining.' in context *c* is not determined by the content in *c* of the constituent sentence 'The sun is shining'. (Lewis 1980 : 95)

That is so because, for Lewis and Stalnaker, the propositional content of 'the sun is shining' in context *c* is the (classical) proposition that the sun is shining at the time and place of *c*. But the proposition expressed by 'Somewhere the sun is shining' is not a function of *that* proposition. 'Somewhere the sun is shining' does not present that proposition (the proposition that it is raining at the place of *c*) as true somewhere ; what it presents as true somewhere is the place-relative proposition that the sun is shining, *punkt*.

Lewis concludes that there is no particular reason for preferring the two-stage picture over the other approach. The only potentially convincing justification that has been offered for singling out 'contents' as a separate level of semantic value for sentences is that they can serve as the objects of propositional attitudes and locutionary acts. But that is a mistake : the contents which the two-step procedure assigns to sentences as semantic values must be compositional, and that means that they must be relativized propositions : functions from rich indices to truth-values. But the contents of propositional attitudes and locutionary acts are, according to Stalnaker, classical propositions : functions from possible worlds to truth-values. This raises a dilemma for the two-stage theorists. Either, like Kaplan, he construes contents as relativized propositions, in which case they can be the semantic values of sentences but they cannot be equated with the objects of attitudes and locutionary acts ; or, like Stalnaker, he construes contents as classical propositions, but then they can no longer serve as semantic values because they are not compositional.

Chapter 14
Rescuing the two-stage picture

Relativizing Stalnaker's propositions

In contrast to Stalnaker's propositions, which are classical, Kaplan's contents are relativized propositions, true with respect to rich indices consisting not only of a possible world (as in Stalnaker's theory) but also of a time and possibly a location. Such relativized propositions could be the semantic values of sentences, by Lewis's standards, since they are compositional : the proposition expressed by 'Somewhere the sun is shining' *is* a function of the place-relative proposition expressed by 'The sun is shining'.

Stalnaker himself considers the possibility of relativizing his propositions.[32] Lewis acknowledges that this would make them compositional, hence a suitable candidate for the status of semantic value, but he has an objection to this potential move :

> Stalnaker does suggest, at one point, that he might put world-time pairs in place of worlds. (…) But this does not go far enough. World and time are not the only shiftable features of context. And also perhaps it goes too far. *If propositions are reconstrued so that they may vary in truth from one time to another, are they still suitable objects for propositional attitudes ?* (Lewis 1980 : 95 ; emphasis mine)

This last objection has a Richardian flavour, and we should be suspicious of it. Have we not seen that relativized propositions *can* be the objects of the attitudes (alongside Austinian propositions) ? Lewis himself, in 'Attitudes *de dicto* and *de se*',

---

[32] « Does a tensed sentence determine a proposition which is sometimes true, sometimes false, or does it express different timeless propositions at different times ? I doubt that a single general answer can be given. » (Stalnaker 1970 : 289)

has argued that the content of attitudes *are* relativized propositions — functions from centered worlds (i.e. world-time-individual triples) to truth-values. He has taken on board, and generalized, Loar's suggestion that what we believe may be 'propositional functions' (Loar), or 'properties' (Lewis), which the believer self-ascribes. In this way Lewis is able to account for the vexing phenomenon of indexical belief — a phenomenon noted by various authors in the second half of the twentieth century, from Geach and Prior in the fifties to Castañeda in the sixties and Perry in the seventies.[33]

Indexical belief

Prior had suggested that, without temporal propositions as possible objects for the attitudes, we cannot account for the attitude of relief expressed by an utterance such as 'Thank Goodness that's over' :

> [This utterance] says something which it is impossible that any use of a tenseless copula with a date should convey. It certainly doesn't mean the same as, e.g. 'Thank goodness the date of the conclusion of that thing is Friday, June 15, 1954', even if it is said then. (Nor, for that matter, does it mean 'Thank goodness the conclusion of that thing is contemporaneous with this utterance'. Why should anyone thank goodness for that ?) (Prior 1959/1976 : 84)[34]

Castañeda and Perry provided a number of similar examples, in which it seems that the object of belief cannot be a classical proposition. Lewis discusses one of Perry's well-known examples : the case of Heimson, the madman who believed he was Hume.

Heimson and Hume both believed to be Hume, but Hume was right and Heimson was wrong. We need something to be the common content of their beliefs. At the same time, the complete content of their beliefs cannot be the same, since one

---

[33] See Chapter 33 for a historical overview.

[34] See Part VII for a discussion of Prior's example, with special reference to Higginbotham's analysis.

belief is true and the other is false. This is one of the cases in which we need to relativize, and to distinguish two levels of content. That is what Lewis does, and his view fits Moderate Relativism perfectly.

Lewis first establishes that the common content of Hume's and Heimson's beliefs cannot be a classical proposition. If it was a classical proposition, it would not be possible for their beliefs to differ in truth-value :

> [Heimson] believes he is Hume. Hume believed that too. Hume was right. If Hume believed he was Hume by believing a [classical] proposition, that proposition was true. Heimson believes just what Hume did. But Hume and Heimson are worldmates. Any proposition true for Hume is likewise true for Heimson. So Heimson, like Hume, believes he is Hume by believing a true proposition. So he's right. But he's not right. He's wrong, because he believes he's Hume and he isn't. (Lewis 1979/1983 : 142)

Lewis concludes by rejecting the assumption that occurs in the antecedent of the conditional : that Hume believed he was Hume *by believing a classical proposition*. Instead, he suggests, Hume believed he was Hume by believing a relativized proposition, true at himself (but false at Heimson), that is, by self-ascribing a property : the property of being Hume. Hume was right, since he actually had the property. Heimson self-ascribed the same property — wrongly, since *he* did not have that property.

There is another option, however : one may *reject the assumption that Hume and Heimson believed the same thing*. This is the option pursued by John Perry, who maintains that the objects of belief are classical propositions.[35] Since classical

---

[35] For Perry, such propositions are believed under varying guises or modes of presentation. So, in the Prior example, Perry would maintain that what is believed is the classical proposition that the time of the conclusion of the painful event is *t*, where the time *t* is thought of under a certain indexical mode of presentation. It is the mode of presentation which accounts for the subject's relief. Or, in the Hume-Heimson case, Perry says that Hume rightly believes the proposition that Hume is Hume, while Heimson wrongly believes the proposition that Heimson is Hume. Each of them believes the relevant proposition under a mode of

propositions must be true or false, *tertium non datur*, if Heimson is wrong and Hume right it follows that they believe distinct propositions, hence it is not true that they believe the same thing.[36] But Lewis has an objection to that move :

> Doubtless it is true in *some* sense that Heimson does not believe what Hume did. But there had better also be a central and important sense in which Heimson and Hume believe alike. (…) Heimson may have got his head into perfect match with Hume's in every way that is at all relevant to what he believes. If nevertheless Heimson and Hume do not believe alike, then *beliefs ain't in the head !* They depend partly on something else, so that if your head is in a certain state and you're Hume you believe one thing, but if your head is in that same state and you're Heimson you believe something else. Not good. The main purpose of assigning objects of attitude is, I take it, to characterize states of the head ; to specify their causal roles with respect to behavior, stimuli, and one another. If the assignment of objects depends partly on something besides the state of the head, it will not serve this purpose. The states it characterizes will not be occupants of the causal roles. (Lewis 1979/1983 : 142-3)

In other words, there is a sense in which Hume and Heimson believe the same thing — this is the 'causal-explanatory' sense, in McGinn's terminology — and there is another sense in which they do not believe the same thing — this is the 'truth-conditional' sense. This distinction corresponds to the duality of belief we talked about in discussing Richard's view. In the causal-explanatory sense, the object of belief must be 'in the head' : it must be purely psychological. In the truth-conditional sense, the object of belief depends upon the context in which the belief is

---

presentation involving his self-concept ; so each believes that *he himself* is Hume, and that captures the sense in which (it seems that) they believe the same thing.

[36] It is not true that they believe the same thing, since they believe distinct propositions ; but it is true that they believe what they believe – distinct propositions – under the same (type of) mode of presentation, corresponding to the indexical sentence they would both utter : 'I am Hume'. (See Chapter 15 on the distinction between 'belief' and 'acceptance'.)

entertained : so it makes a difference to what is believed (in that sense) whether the believer is Heimson or Hume.

For Lewis, the psychological sense is more fundamental :

Mean what you will by « object of an attitude. » But if you mean something that is not determined by the state of the head, and that cannot do the job of characterizing states of the head by their causal roles, then I think you had better introduce something else that can do that job. I would prefer to reserve the term « object of an attitude » for that something else. (Lewis 1979/83 : 143)

Be that as it may, Lewis's claim is that the object of belief, *in the psychological sense*, is not a classical proposition but a relativized proposition, at least when the belief is indexical (as in the Hume/Heimson case).

As for the object of belief in the other, truth-conditional sense, it can easily be captured once we have identified the psychological content of the belief. This is where Moderate Relativism comes in handy. In the truth-conditional sense, the (complete) content of Hume's or Heimson's belief is the Austinian proposition, consisting of the self-ascribed property *and* the situation it concerns, viz. the subject which self-ascribes the property. So Hume's belief and Heimson's belief share the same narrow content (the same *lekton*), but they differ since the other component of the Austinian proposition (the situation component) is Hume in one case, and Heimson in the other. The resulting Austinian propositions have different truth-conditions : one is true iff Hume is Hume, and the other is true iff Heimson is Hume.

Lewis's retractation

If we take Lewis's own theory of *de se* belief seriously, we must acknowledge that his objection to the two-stage picture cannot be maintained. To make Stalnaker's 'middlemen' compositional, hence suitable candidates for the status of semantic value, we may follow Kaplan and construe them as relativized propositions. Contrary to what Lewis suggests, if we do that, we do not lose the Stalnakerian justification that propositions, thus understood, can serve as objects for the attitudes. For it is not true that only classical propositions can be the object of the attitudes. If Lewis (1979)

is right, classical propositions are *never* the object of the attitudes ; or rather, the cases in which they are (or seem to be) can be re-analysed as cases in which the subject self-ascribes the property of inhabiting a world in which the proposition holds.[37] Either way, the propositional attitudes turn out to have relativized propositions as their contents.

Lewis himself noticed that his objection to the two-stage picture collapses, given his own theory of *de se* belief. In preparing for the reprinting of his 1980 article in his *Papers in Philosophical Logic* (Lewis 1998), he added a footnote after the passage I quoted above (p. 00). The passage ends with a rhetorical question :

> If propositions are reconstrued so that they may vary in truth from one time to another, are they still suitable objects for propositional attitudes ? (Lewis 1980 : 95 ; 1998 : 39)

The new footnote contains Lewis's retrospective answer, which sounds like a retractation :

> Yes indeed. For discussion, see my 'Attitudes De Dicto and *De se*'. (Lewis 1998 : 39)

---

[37] Once again, contingency turns out to be a form of indexicality.

Chapter 15
Content, character, and cognitive significance

Three levels, or two ?

Let us take stock. We have seen that beliefs can be individuated in two ways : by their truth-conditions, which often depend upon the context in which they are entertained, or by their intra-individual psychological role. In the latter role, beliefs must be ascribed relativized propositions (properties) as their contents, rather than classical propositions. Since Kaplanian contents are relativized propositions, they can serve as the objects of the attitudes in the psychological sense. This is the epistemic justification for the two-stage picture.

Moderate Relativism as I have spelled it out so far relies on the two-stage picture — that is why it was important to check the acceptability of the epistemic justification. As we have seen, the *lekton* is very similar to Kaplan's 'content'. Both are relativized propositions, and both result from evaluating the indexical expressions with respect to context. So both the *lekton* and the kaplanian content of an utterance are context-dependent ; but they fall short of determining a complete, classical proposition. To get a complete proposition, we need to appeal to context a second time, to fix the relevant index/circumstance with respect to which the *lekton* must be evaluated.

In this framework, which is the kaplanian version of the two-stage picture, three levels are distinguished :

• The (context-independent) meaning of the sentence (a.k.a. Kaplan's 'character') is a function whose argument is the context of utterance and whose value is the contextually variable *lekton* (a.k.a. Kaplan's 'content').

• The *lekton* is a relativized proposition, i.e. a function whose argument is a circumstance/index/situation and whose value is a truth-value.

• The complete content of the utterance is the Austinian proposition consisting of the *lekton* and the contextually relevant situation. It determines a classical proposition, true iff the *lekton* is true at the relevant index.

There is another version of the two-stage picture, however. In contrast to Kaplan, other two-stage theorists do not posit a third, intermediate level of content between the meaning of the sentence and the classical proposition it expresses. Thus Stalnaker's 'middlemen' are classical propositions — functions from possible worlds to truth-values. These classical propositions are the objects of the attitude, according to Stalnaker. The same thing holds for Perry : like Kaplan, he distinguishes two semantic values for utterances, viz. character and content, but he construes contents as classical propositions, and he takes the classical propositions in question to be the object of the attitudes. Perry accounts for indexical belief, not by positing relativized propositions as semantic objects intermediate between characters and classical propositions, but by appealing to characters themselves as the modes of presentation under which classical propositions are believed or disbelieved. So Stalnaker and Perry are both *classical two-stage theorists* : they posit two levels, namely character (sentence meaning) and classical proposition, and eschew the need for an intermediate level of relativized content. In contrast, Moderate Relativism as I have elaborated it so far follows Kaplan in positing such an intermediate level : the *lekton*, distinct both from the context-independent meaning of the sentence and from the classical proposition it expresses in context.

An objection to the three-level version

In terms of this distinction between two versions of the two-stage picture, a devastating objection can be raised against the three-level version endorsed by Moderate Relativism as I have elaborated it. The objection runs as follows : the justification that has been offered for the intermediate level (the *lekton*) is that it can serve as the object of the attitudes in the narrow, psychological sense. But the *lekton* as I have construed it is demonstrably not an appropriate entity for that role, precisely because it is context-dependent. The object of the attitudes, in the narrow psychological sense, must be entirely 'in the head' ; that is the argument that was used to show that the object of the attitudes must be the *lekton* rather than the

classical proposition. But the *lekton* itself does not fit that bill, for it incorporates the contextual values of the indexicals and therefore depends in part upon the context. So if we take the psychological constraint seriously, it is the (context-independent) character, not the *lekton*, which must play the role of object of the attitudes in the narrow sense. Indeed, it is Kaplan's characters, not his contents, that are supposed to play the intra-individual, psychological role of accounting for one's behaviour. As Kaplan himself writes, paraphrasing Perry (1977),

> *We use the manner of presentation, the character, to individuate psychological states, in explaining and predicting action*. It is the manner of presentation, the character and not the thought apprehended, that is tied to human action. When you and I have beliefs under the common character of 'A bear is about to attack me', we behave similarly. We both roll up in a ball and try to be as still as possible. Different thoughts apprehended, same character, same behavior. When you and I both apprehend that I am about to be attacked by a bear, we behave differently. I roll up in a ball, you run to get help. Same thought apprehended, different characters, different behaviors. (Kaplan 1989 : 532, adapted from Perry 1977/1993b : 23)

To recapitulate the objection : If relativized propositions are to be the objects of the attitudes, it is in the psychological sense ; but in that sense there is a much better candidate for the status of psychological content, namely the meaning of the sentence-type. Thus Perry distinguishes two distinct objects for the attitudes, corresponding roughly to the distinction between content and character. We believe a (classical) proposition, he says, by 'accepting' an interpreted sentence. Sentence meanings are the objects of acceptance, and acceptance, in contrast to belief, is purely psychological : it is 'the mind's contribution to belief'.

Belief and acceptance

Remember Watson and Holmes. Holmes says 'The salt is left of the pepper', because the salt is left of the pepper from Holmes's perspective. From Watson's perspective, the pepper is left of the salt; however, Watson is mistaken and he wrongly says 'The salt is left of the pepper'. Holmes and Watson apparently 'say the

same thing' but Holmes is right and Watson wrong. This is like Heimson and Hume ! From the standpoint of Moderate Relativism, one and the same *lekton* (viz. the perspective-relative proposition that the salt is left of the pepper) is expressed, but Holmes and Watson believe different classical propositions, or different Austinian propositions, because they each relativize the *lekton* to their own perspective.

In the classical framework, Barwise claims,

we have nothing in the theory that classifies the similarity in attitudes of Holmes and Watson in cases like these. And it is this similarity that leads them to make the same bodily movements, reaching in the same direction, though toward different objects, when they want the salt. (Barwise 1989a: 240).

But is this true ? Isn't there something, in the classical two-stage theory, which accounts for the similarity in attitudes, i.e. for the narrow psychological content of the common belief Holmes and Watson express ? Of course there is : Holmes and Watson both utter the sentence 'The salt is left of the pepper', and that sentence corresponds to something they both accept. What they have in common, according to Perry, is a certain *belief state* they are both in, and which leads them e.g. to utter 'The salt is left of the pepper' and to reach in a certain direction when they want the salt.

Perry distinguishes the belief relation which we bear to classical propositions, from the acceptance relation which we bear to interpreted sentences when we are in a belief state that would lead a competent user of the language to think or utter that sentence assertively :

Acceptance is not belief ; rather it is an important component of belief. It is the contribution the subject's mind makes to belief. One has a belief by accepting a sentence. Which belief one thereby has also depends upon who the believer is and when the believing takes place — factors that need have no representation in the mind. (…)

In saying that acceptance is the contribution the subject's mind makes to belief, I mean this. When we believe, we do so by being in belief states. These states have typical effects, which we use to classify them. In particular, we classify them by the sentences a competent speaker of the language in

question would be apt to think or utter in certain circumstances when in that state. (Perry 1980/1993b : 53)

So Holmes and Watson share a certain state with a certain causal-explanatory role, and that state corresponds to, and can be classified by, the character of the sentence the state prompts them to utter. Holmes and Watson both 'accept' the interpreted sentence 'The salt is left of the pepper' ; this accounts for the similarity in their behaviour. By accepting this sentence, they believe distinct classical propositions : Holmes believes the proposition that the salt is left of the pepper from Holmes's perspective, while Watson believes the proposition that the salt is left of the pepper from Watson's perspective. Holmes is right, Watson is wrong. Similarly, Hume and Heimson both accept the sentence 'I am David Hume', so they are both in the same belief state, which leads them to act and speak in a certain way. By being in that state, they believe distinct propositions : Hume correctly believes that he, Hume, is Hume, while Heimson incorrectly believes that *he*, Heimson, is Hume.

Back to the single-stage picture

The objection I have raised is very serious, but it does not force us to give up Moderate Relativism. Rather than give up Moderate Relativism, we might simply give up the Barwise-Kaplan interpretation (i.e. the idea that the *lekton* is context-dependent and incorporates the contextual values of indexicals) and go back to Dummett's suggestion that the *lekton* is the meaning of the sentence-type.

Interestingly, this option takes us back to Lewis's single-stage picture. That is so because the objection to the three-level version of the two-stage picture backfires against the classical two-stage theorist who might be tempted to take advantage of it. If the narrow psychological object of the belief expressed by an utterance — that which enables us to capture the causal-explanatory role of the belief — is the meaning of the sentence-type, as Perry suggests, why do we need the two-stage picture in the first place ? We only need to assign to the sentence a certain meaning, which, in context, will determine a truth-value. That meaning may be equated with a complicated function from context-index pairs to truth-values, following Lewis's suggestion. As has already been noticed, the proposition expressed (in the sense of

the classical two-stage theorist) can easily be defined in that framework, so there is no need for the two-stage picture. The single-stage framework will do as well.

In the single-stage framework, Moderate Relativism still makes sense ; indeed it corresponds to Lewis's own position if I understand him correctly. The Austinian proposition expressed by an utterance consists of (i) a context (which Lewis construes as a world-time-individual triple) and (ii) the context-independent *lekton* equated with a function from context-index pairs to truth-values. The utterance is true if and only if the *lekton* is true with respect to the context and the index it determines (the 'index of the context', for Lewis). The *lekton* thus understood is a property of contexts : it is the property a context has iff the relevant function yields truth when evaluated with respect to the context and the index it determines.

In the postscript to 'General Semantics' (Lewis 1983 : 230), Lewis notes that « these 'contexts' are the same as the 'subjects' that self-ascribe properties » in Lewis's article 'Attitudes *De Dicto* and *De se*'. In the Holmes/Watson case, the contexts are world-bound time-slices of Watson and Holmes, who each self-ascribes the property expressed by 'The salt is left of the pepper', namely the property of being a context from the perspective of which the salt is left of the pepper. In the Hume-Heimson case, the contexts are world-bound time-slices of Heimson and Hume, who each self-ascribes the property expressed by 'I am Hume', namely the property of being Hume (or, in the new terminology, the property of being a context whose 'agent' or 'subject' is Hume).

Summing up

The objection I have raised in this chapter establishes that the *lekton*, qua level of content intermediate between context-independent meaning and classical truth-conditions, is not the object of the attitudes in the narrow psychological sense, contrary to what Barwise suggests. So the epistemic justification of the two-stage picture fails when the picture is fleshed out à la Kaplan-Barwise.

At this point there are two options for (Strong) Moderate Relativism. The first option has just been presented: we may give up the Kaplan-Barwise construal in favor of the Dummett-Lewis construal. That is, we may equate the *lekton* with the context-independent meaning of the sentence-type, and the situation with the context

of utterance. On this view, all forms of relativity reduce to indexicality, as Lewis pointed out.

Second option : we may give up the epistemic justification and try to find *another* justification for the Kaplanian, three-level version of the two-stage picture which, following Barwise, we have opted for so far. That is the option I will pursue in the remainder of this book. I will do so by introducing a quite fundamental distinction between *content* and *mode* and showing how it can be dealt with in the SMR framework (Book II). As we shall see in due course (Chapter 30), the mode/content distinction is what ultimately justifies construing the *lekton* as context-dependent.

Book II

EXPERIENCE AND SUBJECTIVITY

Part Five

Content and Mode

Chapter 16

Duality and the fallacy of misplaced information

Sense and force

In Chapter 2, I said that, appearances notwithstanding, Frege accepted Duality (the need for a circumstance of evaluation over and above the content to be evaluated). That is because he drew a distinction between the sense or content of an utterance and its force. An utterance pictures things as being a certain way ; how it pictures things is its sense. But whether the picture is intended as a picture of reality, or as a picture of a fictional state of affairs, is something that cannot be read off the picture itself. This is a pragmatic matter — a matter of how the picture is used.

The assertive force which characterizes serious use (as opposed to fictional or ironical use) can be cashed out in terms of a certain relation of the utterance to the actual world. The speaker who makes an assertion describes the actual world as being thus and so. The speaker who, in Grice's terms, 'makes as if to say' does not stand in that relation to the actual world. Frege, of course, would not say that in such a case the speaker stands in that relation to some other, unactualized world (or worlds). He claims that the author of a fiction offers a picture which may be beautiful or interesting but which is *not* a picture of anything. *No* relation is established to something against which the picture could be evaluated. Fiction, Frege says, is sense without reference ; and without reference, there is no possibility of truth. (Truth is a special case of reference, for Frege.)

Assertive force is thus a matter of the utterance's standing in an appropriate relation to the world, a relation in virtue of which the sense of the utterance can be evaluated as incorrectly or correctly describing the world. Besides this distinction between sense and force, Frege also drew a distinction between the complete sense of an indexical utterance and the partial sense which the uttered sentence, by itself, carries. The sense of an indexical sentence is incomplete, Frege says : it must be completed by the circumstances of the utterance. Only the resulting 'hybrid symbol' (sentence + context) carries a complete sense. The important point is that the force

of the utterance comes *in addition to* that complete sense, and it cannot be made an aspect of it — not even through the addition of a prefix like 'I assert that', or 'it is true that'. If the speaker is not serious, adding extra words and making the sense more complex will not change the situation.

On this admittedly unconventional interpretation, Frege's view is compatible with Moderate Relativism. The complete import of an utterance involves a situation (the actual world) and a content to be evaluated with respect to that situation (the complete sense). That content is context-dependent, and it is classical (i.e. fully propositional, by Fregean standards). This corresponds to one of the versions of Moderate Relativism I said we would have to adjudicate between. In order to adjudicate beween them, the crucial issue we have to address is that of the status of the *lekton* : is it context-dependent, and if it is, is it classical or relativized ? However, I will postpone discussion of that issue until Book III. For the time being, I want to investigate the force/content distinction, and its counterpart in the mental realm, in some detail since it can be seen as a special case of the 'Austinian' duality we are concerned with in this book.

Misplaced information

Illocutionary force, in the SMR framework, is handled through the specification of the situation against which the utterance is meant to be evaluated. Just as an assertion that *p* presents the fact that *p* as holding in the actual world, a command that *p* presents the fact that *p* as holding in a situation the addressee is asked to actualize, that is, to bring about. A promise that *p* is similar, with one difference : the actualizer is supposed to be the speaker, not the addressee.

Commands and promises provide a useful illustration of a common mistake, which consists in ascribing to the propositional content of a speech act a property that is actually a feature of its illocutionary force. Thus Searle famously claimed that speech act types are characterized by a number of properties, including properties of their propositional contents. The content of an order must be a future act of the addressee, he said. If I order John to wash the dishes, the content of the order is the proposition that John will wash the dishes — the same proposition that would be the content of an assertion if I *said that* John will wash the dishes, or of a question if I *asked whether* John will wash the dishes. While assertions or questions can have

any proposition as their contents, order and directive speech acts more generally are supposed to be more constrained : *their content can only be a proposition bearing upon a future act of the hearer*, as in this particular case. Likewise, promises are such that their contents can only be a future act of the speaker (Searle 1969 : 57-58). Now Searle's claim rests on what Barwise and Perry call the 'fallacy of misplaced information' (Barwise and Perry 1983 : 38). As I wrote in *Meaning and Force*,

> The notion of the hearer's (or speaker's) future behavior does have a role to play in the analysis of commissives and directives, but… the proper place for this notion is in the analysis of the *illocutionary force* of directives and commissives, not in the analysis of their propositional content. Any proposition whatever can be the content of a directive or a commissive ; it suffices that the speaker's utterance express his intention that the hearer (or the speaker), by virtue of this utterance expressing this intention, behave in such a way as to make the proposition true. The 'hearer's (or speaker's) future behavior' consists in bringing about the state of affairs that is the content of the speech act. It is not an intrinsic aspect of that state of affairs. (Recanati 1987 : 163, originally in Recanati 1981: 189 ; see also Sperber 1982 : 47)

As I pointed out in the same book, Searle's view has unfortunate consequences : it forces the theorist to treat as indirect the illocutionary act of promise performed by means of a sentence such as 'You will win' (said by a tennis player to his opponent), while the same act would be performed directly if the speaker had uttered 'I will lose' instead. This sounds arbitrary. Again, any future state of affairs can be the content of an order (or a promise) if the utterance manifests the speaker's intention that, because of this utterance, the addressee (or the speaker) will bring about that state of affairs. Even the futurity of the state of affairs is not a genuine constraint on the propositional content of commissives or directives. To be sure, it is a natural constraint that humans can only actualize the future, since the past is beyond their control ; but that is not a conventional, linguistic constraint.

In the SMR framework the content of the illocutionary act corresponds to the *lekton,* and its force corresponds to the type of situation with respect to which the *lekton* is meant to be evaluated. Thus an order that *p* represents the fact that *p* (content) and it represents it as holding in a situation *brought about by the hearer as*

a causal result of the speech act (via a causal chain involving the hearer's recognition of the speaker's illocutionary intention to issue an order). A promise that *p* similarly represents the fact that *p* as holding in a situation brought about by the speaker as a causal result of the speech act. In both cases, the notion of a future act of the speaker or the hearer comes into the picture through the situation component of the Austinian proposition which is the complete content of the speech act.

From illocutionary force to psychological mode

In his later work, Searle emphasized the similarity between speech acts and mental states. Both have a propositional content, and both have a dual structure, with the 'psychological mode' corresponding, on the side of mental states, to the illocutionary force on the side of speech acts. To assert that *p* is to perform a speech act whose content is the proposition that *p* and whose force is that of an assertion. Likewise, to believe that *p* is to be in an intentional state whose content is the proposition that *p* and whose psychological mode is that of belief. Now, because of that dual aspect, it is all too easy to 'misplace' some of the information carried by a mental state, by ascribing to the content of the state a piece of information that is actually carried by its mode — just as Searle misleadingly ascribed to the propositional content of directive and commissive speech acts what was actually a feature of their illocutionary force.

As we shall see in Chapter 17, Searle himself has been guilty of such misplacement in his insightful analysis of perceptual experiences (a sub-class of Intentional states). Besides the distinction I have just introduced, between the content of the state and its mode, Searle draws a second distinction, between two components in the overall truth-conditions of perceptual experiences. He distinguishes two conditions such that the experience counts as satisfied (veridical) if and only if they are both met : the *primary condition* and the *self-referential condition*. (Those are my terms, not Searle's.) For him, both the primary condition and the self-referential condition are determined by the propositional content of the experience ; he therefore construes that propositional content as self-referential. If I am right, however, the propositional content of the experience only determines the first component of its overall truth-conditions — what I called the primary condition. The self-referential condition is not determined by the propositional content of the state,

but by its mode. It follows that the propositional content of a perceptual state is not self-referential, even if its overall truth-conditions are.

Chapter 17

The content of perceptual judgments

Searle's analysis

Visual experiences have propositional content, Searle claims. The subject does not merely see a flower, he sees a flower there, and this, fully spelled out, means that he sees *that there is* a flower there. I understand Searle's contention as follows : the subject, as part of his visual experience, makes certain perceptual judgments, which determine the (conceptual) content of the experience. So the propositional content of the experience is the content of the judgments that are immediately (i.e. noninferentially) based on it.[38]

Let us ask what the truth-conditions of the perceptual judgement are. For the perceptual experience to be veridical, Searle says, there must actually be a flower there, but that is not sufficient. In addition to the primary condition that there be a flower there, an extra condition must be met : it must be the case that the presence of a flower there causes the visual experience (including the judgment that is part and parcel of the experience). Insofar as the propositional content of the state is what determines its truth-conditions, the content of the visual experience turns out to be more complex than one might have thought. The content of the visual experience is not the simple proposition *that there is a flower there*, but the conjunctive proposition *that there is a flower there and there being a flower there causes this visual experience* (where 'this visual experience' reflexively refers to the experience of which this is the propositional content). Searle therefore provides the following analysis of the perceptual experience of seeing a flower there. As in the analysis of speech acts, the stuff in capital letters indicates the mode, while the material within the parentheses specifies the content.

---

[38] I know this is not how Searle himself puts the matter, but the difference in formulation is not important in the context of the present discussion.

VIS EXP (that there is a flower there and that there is a flower there is causing this visual experience)

According to this analysis, the intentional state is a visual experience (mode) and its content is a conjunctive proposition whose second conjunct refers to the very experience of which this is the content, while the first conjunct specifies 'the state of affairs perceived' (Searle 1983 : 48 ; I suppose this means something like : the state of affairs which the speaker takes himself to be perceiving). The first conjunct determines a proper part of the judgment's truth-conditions, namely the primary condition : that there be a flower there. But there is another truth-condition, determined by the second conjunct. That is the self-referential condition : that the perceptual experience be caused by 'the rest of its truth-conditions', that is, by the state of affairs whose existence is the primary condition for the judgment to be veridical. Because of the second condition, the overall truth-conditions of the perceptual judgment are self-referential : for the perceptual judgment to be veridical, the state of affairs which the speaker takes himself to be perceiving must actually exist (primary condition), and it must be what causes the speaker's perceptual judgement that that state of affairs exists (self-referential condition).

Introducing the mode/content distinction

I fully agree with Searle that there are these two components in the truth-conditions of perceptual judgements. But I deplore Searle's claim that the overall truth-conditional content of the judgment, with its two components, is the propositional content of the visual experience (as distinguished from its mode). Insofar as I can tell, this is the fallacy of misplaced information once again.

That the state of affairs represented (there being a flower there) causes the representation of that state of affairs is a condition that has to be met for the representation in question to count as a *perception* (rather than, say, an expectation). It follows that the self-referential condition is determined by the perceptual mode of the state, not by its content.[39] For a representation that *p* to count as a perception

---

[39] Thus I concur with Kent Bach : « It is not the content of a visual experience that determines that its cause is its object. Rather, it is the psychological mode. It is the fact that the

that *p*, it must be the case that the representation is caused by the fact that *p* ; but what is represented is only the fact that *p*. In other words : the content of the state (viz. the proposition that *p*) only determines the primary condition ; the perceptual nature of the state is what determines the self-referential condition. Together, the content and the mode determine the overall truth-conditions of the state.

Searle might insist that, for him, the propositional content of a state just *is* what determines its truth-conditions. Consider the following passage, where he introduces the idea of propositional content :

> Every Intentional state consists of an *Intentional content* in a *psychological mode*. Where that content is a whole proposition and where there is a direction of fit, the Intentional content determines the *conditions of satisfaction*. Conditions of satisfaction are those conditions which, as determined by the Intentional content, must obtain if the state is to be satisfied. For this reason the *specification* of the content already is a *specification* of the conditions of satisfaction. Thus, if I have a belief that it is raining, the content of my belief is : that it is raining. And the conditions of satisfaction are : that it is raining.
> (Searle 1983 : 12-13)

Given this characterization of content, it is legitimate to use the overall truth-conditions of perceptual experiences as evidence to determine what their content is. Following this procedure, we end up with Searle's conclusion : the propositional content of a visual experience is self-referential.

But I do not think we can accept this defense. Searle cannot simply 'define' the propositional content of a state as that which determines its truth-conditions. Or rather, he can, but then the claim that the content so defined is (to some extent) independent of the mode becomes an empirical claim, and it is that claim which I reject.[40] There is a sense of 'content' in which the content of perceptual judgments is

---

experience is perceptual that determines that its object(s) is that which causes, in the appropriate way, that very state. » (Bach forthcoming)

[40] In response, Searle might point out that he never claimed that the content was independent of the other component (mode or force) : the very idea of a 'propositional content condition' argues against such independence. This point is well-taken, but I reply that the distinction

self-referential ; but the 'content', in *that* sense, is determined in part by the mode, and in part by… the content of the state in a *different* sense, namely, the sense in which 'content' and 'mode' are two independent dimensions which together constitute the state. My point, therefore, is that Searle cannot simultaneously maintain that the content is what determines the conditions of satisfaction, and that the content is independent of the mode. There are two distinct notions of content : one that is involved in the mode/content distinction inspired from speech act theory, and another one that is involved in the claim that the content of a state is what determines its (overall) truth-conditions.

A second line of response is available to Searle. He may insist that *it is part of the subject's perceptual experience* that that experience is caused by its object. The 'causation' component is not external to the content of the experience, but an integral part of it. With this I agree — but I do not think Searle's position is thereby justified. The subject is *aware* of the perceptual nature of his experience : he knows he is perceiving rather than, say, expecting or imagining. So there is a sense in which the complete content of his experience is self-referential : but the 'complete content' of the experience involves more than the propositional content – it also involves the psychological mode, of which the subject is aware and which determines the additional, self-referential condition.

The proper analysis

Once again, the analogy with speech acts can be illuminating. There is a self-referential element in speech acts too, as several authors have pointed out. For example, an order represents a certain state of affairs as a state of affairs such that the addressee complies with the order only if he brings it about. But bringing about

---

between 'content' and 'mode' (or 'force'), by itself, implies that these are relatively independent dimensions. However lightly we construe this 'independence' condition, it will not be satisfied (I claim) if we follow Searle and *define* content as that which determines satisfaction conditions. Searle himself concurrently uses a distinct notion of content, namely (for any type of speech act or Intentional state F) *what* the subject Fs : what she believes, what she asserts, what she promises, what she perceives, etc. The content of the act/state, in this sense, satisfies the independence condition, and it is not self-referential.

the state of affairs in question is not sufficient for compliance : it must be the case that the addressee brings about the state of affairs *as a result of being ordered to do so.* An order, therefore, presents a state of affairs as something which the hearer is to bring about as a result of this order. So there are two components in the obedience conditions of an order : the hearer must do something (primary condition), and the order must be what causes the hearer to do that thing (self-referential condition). But it would be a mistake to consider the second, self-referential condition (that the order is what causes the hearer to do *x*) as determined by the *content* of the order (i.e. what is ordered), as opposed to its force. The speaker orders the hearer to do *x*, period ; he does not order her *to do x because of this order*. Yet the fact that the utterance *is* an order (rather than an assertion) means that the hearer's doing *x* will count as satisfying the speech act only if the hearer's doing *x* is caused by the speech act in the proper way. This is a feature of the illocutionary force of ordering. So the correct representation of the speech act of, say, *ordering the hearer to wash the dishes* is not

ORDER (that the addressee washes the dishes and that the addressee washes the dishes is caused by this order)

but simply

ORDER (that the addressee washes the dishes)[41]

I think exactly the same considerations apply in the perception case. The proper representation of the perceptual experience of the subject who sees a flower there is not

VIS EXP (that there is a flower there and that there is a flower there is causing this visual experience)

---

[41] As Savas Tsohatzidis pointed out to me, the analysis of directives I am criticizing here as incorrect (that which construes their content as self-referential) is explicitly endorsed by Searle in *Intentionality* (Searle 1983: 86).

but simply

VIS EXP (that there is a flower there)

In other words, the content of the perception of a flower is the fact that there is a flower there. That fact can be represented in all sorts of modes ; for it to be represented in the perceptual mode, it must be the case that the fact itself causes the representation. But this feature, hence the self-referential component whose importance Searle rightly emphasizes, is a property of the perceptual mode of representation, not a property of the content of perceptual representations. (Or, if it is a property of their content, it is not in the sense in which 'content' contrasts with 'mode', but in the sense of 'overall' or 'complete' content discussed in Chapter 3.)

[It is one of the advantages of the SMR framework that the distinction between content and force, or between content and mode, can be clearly articulated. The speaker who orders the addressee to wash the dishes represents the fact that the hearer is going to wash the dishes, and represents it as holding in a very specific situation, namely a situation resulting causally from the speech act (via a causal chain involving the addressee's recognizing the speaker's intention, etc.). Similarly, the speaker who sees a flower there entertains a representation whose content is the proposition that there is a flower there ; but that proposition, being entertained in the perceptual mode, is meant to be evaluated with respect to a very specific situation, namely the subject's *perceptual situation* : a situation which the subject is causally affected by through his senses and which, in particular, causes the occurrence of the mental representation in question.]

Chapter 18

Episodic memory

What has been said about the content of perceptual judgments naturally extends to the content of (episodic) memory. But memory raises special difficulties, and even though the conclusion we will reach is very much the same as that we reached in the analysis of perception, it is worth taking a fresh start and considering what is specific (and specifically problematic) about memory.

The conjunctive analysis

Episodic memories are mental states which presuppose other mental states, namely perceptual experiences, to which they are related both causally (the memory derives from the perceptual experience, which leaves it as a 'trace') and semantically (the memory inherits the content of the perceptual experience). As Evans puts it, we need to be able to use information gathered through perception at a later time, and for that we need a mechanism for retaining information. Memory is that mechanism. Or rather, it is a family of mechanisms, corresponding to distinct ways of retaining information. In *semantic memory*, we retain the beliefs once formed on the basis of perception (or on any other basis — what is retained is only the output of the belief-fixation mechanism, so the etiology of the belief is irrelevant). In *episodic memory*, we retain the perceptual experiences themselves.[42] I will be concerned only with the second type of memory, where a memory state is an experience similar to the perceptual experience from which it derives.

The crucial feature of episodic memories, besides their being experiential states, is that they share their content with the perceptual experiences from which

---

[42] « Episodic memory is memory for personally experienced events… It transcends semantic memory by being ego-centred : its contents include a reference to the self in subjective space/time » (Tulving 1993 : 67). On the implicit self-reference at work in episodic memory, see Part VI.

they derive : they represent the same event or scene as the perceptual experience which is their ancestor. That is what episodic memory is for — it is supposed to *replicate* the perceptual experience. However, that is not all there is to say about memory. Even though the function of memory is to replicate the perceptual experience and, in particular, to carry the same content as that of the perceptual experience, still there is a fundamental difference in content between memory and the perception on which it is based.

I see that there is a large tree standing in front of me. Later, I remember that there was a large tree standing in front of me. This way of putting things reveals the essential difference : perception and memory relate to the same scene or event, but in memory the scene or event is *presented as past*. I remember that there *was* a large tree standing in front of me.

There is an obvious tension between the two elements I have just mentioned. Does a memory state have the same content as the perceptual state from which it derives, or does it have a different content ? The answer seems to be : both ! The scene or event the memory is about is clearly the same as the scene or event the initiating perception is about (that is what makes memory memory), but it is not the same since perception represents a present event (and represents it *as present*) while memory represents a past event (and represents it *as past*).

Faced with this tension, the obvious move is to distinguish two components in the content of a memory. One component is common to the memory and the perception from which it derives. Since that component is shared, the memory 'has the same content as' the perception. But the content of the memory involves an additional element, which is responsible for the difference between the memory and the perceptual state. Thus William James defines memory as 'the knowledge of an event…, of which in the meantime we have not been thinking, *with the additional consciousness that we have thought or experienced it before*' (James 1890 : 648, cited in Hoerl 2001 : 326 ; my emphasis). Here the 'knowledge of the event' is what is common to the perceptual state and the memory state ;[43] the additional component

---

[43] « [Suppose] I have on some occasion, a strong recollection of a scene at which I was recently present, and no specific reason to doubt it. (…) My recollection may well be wrong ; but, so long as I trust it, I cannot separate the knowledge I suppose myself to have now from the knowledge I surely had at the past time. For the former is derived from the latter ; more

is responsible for the 'feeling of pastness' which differentiates memory from perception.[44]

On this view the content of memory is conjunctive. The first conjunct is what is shared with perception ; the second conjunct differentiates memory from perception. The first conjunct represents a scene or event in the world ; the second conjunct represents the subject's past perceptual experience of that scene or event — hence the second conjunct is meta-representational. This is very similar to Searle's analysis of the content of perceptual experiences. According to Searle, the visual experience of a flower has conjunctive content : *there is a flower there and that there is a flower there causes this visual experience*. The first conjunct is about the flower, the second conjunct is about the subject's experience of the flower. Since the experience mentioned in the second conjunct is the very experience whose content is being analysed, the second conjunct makes the content of perception self-referential. Now, the same idea applies to memory, according to Searle. He analyses the content of episodic memory as follows :

> The memory of seeing the flower represents both the visual experience and the flower and is self-referential in the sense that, unless the memory was caused by the visual experience which in turn was caused by the presence of (and features of) the flower, I didn't really remember seeing the flower. (Searle 1983 : 95)

The memory, for Searle, represents the flower, the past visual experience as caused by the flower, and the present memory as caused by the past visual experience of the flower. The flower is represented and at the same time the representation itself is represented as a memory, that is, as causally deriving from an earlier, perceptual representation of the flower.

---

exactly, it simply *is* the knowledge I had as an eyewitness, maintained in being. » (Dummett 1993 : 414-415)

[44] That feeling is « a distinctive, unique awareness of re-experiencing here and now something that happened before, at another time and in another place. The awareness and its feeling-tone is intimately familiar to every normal human being » (Tulving 1993 : 68).

The metarepresentational analysis

There is a problem with the conjunctive analysis. It takes the content of the memory experience to consist of the content of the perceptual experience together with an additional element. But in what sense does the content of the memory experience contain that of the perceptual experience ? In perceiving the flower, I judge that there is a flower there. In remembering the flower, I do *not* judge that there is a flower there — only that there was one. So it does not seem that one can get to the content of memory by simply adding something to the content of perception.

In response to this objection, we may revise the analysis and interpret differently the idea that the content of memory contains the content of perception as a proper part. We may give up the conjunctive analysis and consider the content of the perceptual experience as occurring in the content of the memory not as an independent conjunct, but as a subordinate, embedded part. Instead of '*p* and I once perceived that *p*' the content would be simply 'I once perceived that *p*'. (To this, following Searle, we may add the further idea that the present memory state is caused by that earlier perception : 'I once perceived that *p* and that perception is causally responsible for this memory experience'.) The content of the perception occurs here as *mentioned* in the metarepresentational component, which now exhausts the content of the memory instead of being only a part of it. This view I call the metarepresentational analysis.

For this analysis to work, a few adjustments are necessary. First, the paraphrase 'I once perceived that *p*' is misleading. It does not capture the experiential component of memory — the fact that in memory, the nonconceptual content of the perceptual experience is retained (to some extent at least). A better paraphrase would be : 'I had perceptual experience XXX', where 'XXX' does not merely specify the conceptual content of the perceptual judgment, but, as it were, 'quotes' the perceptual experience directly. Another adjustment is needed to capture the epistemic value of memory, namely the fact that, if I remember that *p*, I am thereby justified in judging that it was the case that *p*. That would not be the case if the content of memory was simply a representation of a (possibly non veridical) past perceptual experience. For that reason, Jordi Fernandez has suggested adding the condition that the past perceptual experience which memory represents is *represented as veridical* (Fernandez forthcoming).

The problem with the metarepresentational theory is that it presents memory as primarily about our perceptual experience, and only indirectly about the world. I share Evans's protest that

> We no more have, in memory, information that is primarily about our past experiences than we have, in perception, information which is primarily about our present experiences. Just as perception must be regarded as a capacity for gaining information about the world, so memory must be regarded as a capacity for retaining information about the world. (Evans 1982 : 240)

In this respect the conjunctive theory fares better than the metarepresentational theory. The conjunctive theory integrates a metarepresentational component (corresponding to the subject's consciousness of being in a state which causally derives from a previous perception), but it also incorporates the direct representation of a state of affairs in the world, and straightforwardly captures the idea that the memory state retains the content of the perceptual state.

The problem we have raised for the conjunctive theory was that memory does not preserve the content of perception but transforms it by putting it, as it were, in the past tense. That is why it does not seem that one can get to the content of memory by simply adding something to the content of perception. Adding is not enough ; one needs to subtract something as well — namely the present tense or, less metaphorically, the feeling of presentness which colors the perception of the scene. In shifting from perception to memory, this feeling is removed, and it is replaced by the feeling of pastness which colors the representation of the scene in memory.

In the moderate relativist framework, this problem simply does not arise. As in the conjunctive analysis, there are two components, one of which represents a state of affairs in the world. But in the SMR framework, that component is temporally neutral, so there is no objection to saying that it is common to perception and the resulting memory.

The complete content of a perceptual state is analysed into (i) a content in the narrow sense (the *lekton*) and (ii) a situation with respect to which that content is

supposed to be evaluated. The complete content is distributed, and that means that what the situation component supplies need not be replicated in the *lekton*. Now the content of a perceptual experience is relative to the situation of perception. This relativity extends to time : the content of perception is temporally neutral, but it is evaluated with respect to the time of the perceptual experience. So the subject has, at *t*, a perceptual experience the content of which is the temporal proposition that there is a flower there, and that proposition is presented as true at *t*, the time of the present perceptual experience.

Since the content of the perception, in the narrow sense, is temporally neutral, there is no objection to saying that it is preserved in memory. In memory, the same temporally-neutral proposition that there is a flower there is presented as true with respect to the situation (and the time) of the *earlier* perceptual experience rather than the situation (and the time) of the *present* memory experience. In the analysis of memory just as in the analysis of perception, the temporal element is carried by the situation of evaluation.

Phenomenology and the mode/content distinction

At this point the metarepresentational theorist can make the same inadequate response which, in chapter 17, I said Searle could make. It runs as follows : the feeling of pastness, just like the feeling of presentness that accompanies perception, is an aspect of the *content* of memory. There is a clear *phenomenological* difference between memory and perception, having to do with their respective temporal orientations. This cannot be expelled out of the content, however narrowly we construe that notion of content. So, for example, Mike Martin writes:

> The mere possibility that recall and experience might coincide in content
> raise[s] the worry that the phenomenology of recall might *then* have to be
> identified with that of sensory experience. (Martin 2001 : 278 ; emphasis mine)

But there is absolutely no reason to consider that phenomenology supervenes on content in the *narrow* sense. The mode contributes to the phenomenology as much as the content, since the mode is something the subject is aware of. In the memory mode, the content is presented as true with respect to a past perceptual

situation, hence the scene represented is felt as past. In the perception mode, the content is presented as true with respect to the current perceptual situation, hence the scene represented is felt as present. This introduces a difference in the complete content of the respective states, a difference which the phenomenology reflects.[45]

---

[45] We see, once again, that the content of a state, in the narrow sense, does *not* capture its cognitive significance, contrary to what Barwise suggests. For cognitive significance and phenomenology are inseparable, and it we treat the mode (hence the situational component it determines) as relevant to the phenomenology, it will have to be relevant to cognitive significance as well. See Chapter 40 for a related point.

Part Six

Immunity to error through misidentification

Chapter 19
Implicit self-reference

The egocentricity of perceptual judgments

I see that there are clouds. The content of my visual experience is very simple : it is the fact (if it is a fact) that there are clouds. But of course the complete content of my perception is richer than that. As Searle points out, my perception is not veridical if and only if there are clouds ; it is veridical if and only if *there are clouds at the place and time of perception, and there being clouds at the place and time of perception is what causes (in the right way) my perceptual judgment that there are clouds*. To get the complete content of the perceptual judgment, therefore, we need to add something to its content in the narrow sense, namely some kind of reference to the place and time of perception and to the causal relation between the perception and the fact perceived.

To say that those extra pieces of information are not part of the content of the representation, in the narrow sense, is to say that they are not represented in the way in which the clouds are represented. They are not explicitly represented, but find their way into the complete content of the representation because of the *functional role* of that representation, determined by its mode. Qua perception, the representation is caused by the situation around the subject, and it serves to guide the subject's action in that situation. The representation, therefore, is *relative to* the subject's situation, and it is true if and only if the fact represented (viz. that there are clouds) holds in that situation. But the subject and his situation are not themselves elements of the fact represented ; they are not constituents of the *lekton*.

In 'Perception, Action, and the Structure of Believing', John Perry writes :

> The information that we get at a certain spot in the world is information about objects in the neighborhood of that spot in a form suitable for the person in that spot. As long as this is the only source of information we have about

ourselves, we need no way of designating ourselves, indexical or insensitive. Our entire perceptual and doxastic structure provides us with a way of believing about ourselves, without any expression for ourselves. (Perry 1986a/1993b : 148-149)

Perry often expresses the same idea by saying that the subject is an 'unarticulated constituent' of the perceptual judgment. The subject's perceptual judgment is about *his* perceptual situation, but this fact is guaranteed by the role of his representation, qua perceptual judgment. The subject himself does not need to be explicitly represented, since the representation can only be about him and his situation.

The internal mode

An obvious area of application for Perry's idea of the self as unarticulated constituent of certain judgments is the phenomenon of *immunity to error through misidentification*, discussed by Wittgenstein (1958), Shoemaker (1968), and Evans (1982), *inter alia*. If, on the basis of proprioceptive/kinaesthetic experience, I judge that my legs are crossed, I cannot be mistaken as to the person to whom I ascribe the property of having one's legs crossed. It is not possible that I correctly judge that someone's legs are crossed, and misidentify the person whose legs are crossed — thinking that it is myself, while in fact it is someone else. Note that this sort of mistake can easily occur if the judgment is not made 'from inside' but is, say, based upon the perception of my body in the mirror. If I see that my legs are crossed, I may well misidentify the person whose legs are crossed : what I take to be *my* legs in the mirror may be someone else's. When the belief that my legs are crossed is gained through experiencing my own body from inside, no such mistake is possible.

The explanation of this phenomenon is straightforward, and it involves the distinction between mode and content once again. Information gained on the *proprioceptive/kinaesthetic mode* — henceforth to be called the *internal mode* — can only be about one's own body. This is a contingent fact, but it is a fact nonetheless. So the (narrow) content of the bodily experience is not a complete subject-predicate proposition involving a certain person (myself) and a property (being cross-legged or

having one's legs crossed). The content is a relativized proposition, true at a person, and the internal mode determines the person relative to which that relativized content is evaluated : myself. In Perry's terms, the self is an unarticulated constituent of the judgment. The content of the judgment, therefore, is not the proposition that *my* legs are crossed, strictly speaking. What is articulated is only the person-relative proposition that one's legs are crossed. Since I am not explicitly 'identified' as the person whose legs are crossed, I cannot be mis-identified. The judgment, Evans says, is 'identification-free' (Evans 1982 : 180-81). The identity of the person of whom the property holds is secured by the architecture of the system, which guarantees that 'internal' information can only concern the subject's own body.

Implicit and explicit self-ascriptions

I contrasted the case in which I know 'from inside' that my legs are crossed, and the case in which I see that in the mirror. In the latter case, the person to whom the property of having one's legs crossed is attributed is explicitly represented : it is 'articulated' or 'identified' in the content of the representation, which is conjunctive (*Someone's legs are crossed, and I am that person* — the second conjunct is what Evans calls the 'identification component'). In the former case, that person is implicitly determined by the mode, but it is not represented, hence the statement is 'identification-free' : no mistake can be made regarding specifically the person to whom the property is applied. There is no identification component in the content of the judgment. (See Millikan 2004 : 179 for an analogous distinction between 'ego-implicit' and 'ego-explicit' representations.)

Despite this difference, it is important to realize that the self is an unarticulated constituent in the mirror case also. When, looking at the mirror, I see my legs crossed and judge that they are, the perception is about myself in two senses. In a first sense, the perception is (explicitly) about myself because it is *my* legs which, on the basis of my visual experience, I judge to be crossed. Since I am explicitly identified as the person whose legs are crossed, the judgment is susceptible to errors through misidentification. But the perception is (implicitly) about myself also in the sense that *I am the one who sees* the mirror and what it shows. That the perception

is mine and concerns my surroundings rather than someone else's, is something which is guaranteed by the architecture of the system hence gives rise to immunity. So, in the mirror case, the perceptual representation I entertain is such that one property (the property of *seeing in the mirror that my legs are crossed*) is ascribed to myself implicitly, in virtue of the perceptual mode of representation, while another property (the property of *having one's legs crossed*) is ascribed to myself explicitly, in virtue of the content of the perceptual judgment. In the internal case, I self-ascribe *both* properties implicitly, in virtue of the internal mode of representation. What is represented is only the property of having one's legs crossed : its ascription to myself is not an aspect of the content of the representation (in the narrow sense), but something that follows from the mode of representation.

Chapter 20

Weak and strong immunity

Shoemaker on circumstantial vs absolute immunity

In 'Self-Reference and Self-Awareness', Shoemaker draws a distinction between two forms of immunity to error through misidentification, one of which is stronger than the other. A first person statement like 'I see a canary' has *absolute* immunity, Shoemaker claims, because the subject cannot be mistaken as to who is seeing :

> I can be mistaken in saying 'I see a canary', since I can be mistaken in thinking that what I see is a canary or (in the case of hallucination) that there is anything at all that I see, but it cannot happen that I am mistaken in saying this because I have misidentified as myself the person I know to see a canary. (Shoemaker 1968/2003 : 8)

The situation is different, he says, with first person statements like 'I am facing a table', which only have *circumstantial* immunity :

> A statement like 'I am facing a table' does not have this sort of immunity, for we can imagine circumstances in which someone might make this statement on the basis of having misidentified someone else (e.g., the person he sees in a mirror) as himself. But there will be no possibility of such misidentification if one makes this statement on the basis of seeing a table in front of one in the ordinary way (without mirrors, etc.) ; let us say that when made in this way the statement has 'circumstantial immunity' to error through misidentification relative to 'I'. (Shoemaker 1968/2003 : 8)

I think Shoemaker is confused here. Immunity is *always* circumstantial, that is, always relative to a 'way of gaining information' (Evans). Just as I can acquire the information that I am facing a table in a way that does not give rise to immunity, I can acquire the information that I see a canary in such a way too. For example, imagine I am a world-famous neuroscientist suffering from blindsight and undergoing a typical blindsight experiment. Images are shown to me, which I am supposed to identify. Presently an image of a canary is displayed but, being a blindsight patient, I am not aware of seeing a canary (even though I can be shown to have — subpersonally — identified the canary). At the same time, however, the electric activity of some of my neurons involved in visual identification is being recorded and amplified online in the form of a crackling sound, which I can hear. Thanks to an elaborate theory of mine, I am able (or think I am able) to identify what I see from the sound the neurons make. In such a situation I may assert 'I see a canary' because I hear what I take to be the typical sound of canary-neurons. My assertion is clearly not immune to misidentification : for, unbeknown to me, I might well be listening to the neurons of some other patient undergoing experiment in the next room.

I conclude that the statement 'I see a canary' is immune to error through misidentification only if the information that I see a canary is gained in the normal way — through ordinary perceptual awareness. In such a case, the content of the perceptual experience is the proposition that there is a canary there, and the subject's involvement — his or her relation to the scene perceived — is implied by the perceptual mode. In the blindsight case, in contrast, the subject's judgment that he sees a canary has a complex, conjunctive content : the subject judges that someone is seeing a canary (as indicated by the unmistakable sound of canary-neurons) and that he is that person (since he takes himself to be the person whose neurons he is listening to). Because of the second conjunct, the statement is identification-dependent, hence not immune to error through misidentification.

Two types of (implicit) self-ascription

Even though the distinction between absolute and circumstantial immunity is mistaken, there is a legitimate distinction between two kinds of immunity to error through misidentification, which I call 'grade-1' and 'grade-2' immunity. They correspond to two types of implicit self-ascription which are not usually distinguished, but should be.

Consider a well-known passage from *the Bounds of Sense*, cited by Shoemaker and Evans, where Peter Strawson discusses the phenomenon of immunity :

> When a man (a subject of experience) ascribes a current or directly remembered state of consciousness to himself, no use whatever of any criteria of personal identity is required to justify his use of the pronoun 'I' to refer to the suject of that experience. It would make no sense to think or say : *This* inner experience is occurring, but is it occurring to *me* ? (This feeling is anger ; but is it I who am feeling it ?) Again, it would make no sense to think or say : I distinctively remember *that* inner experience occurring, but did it occur to me ? (I remember that terrible feeling of loss ; but was it I who felt it ?) There is nothing that one can thus encounter or recall in the field of inner experience such that there can be any question of one's applying criteria of subject-identity to determine whether the encountered or recalled experience belongs to oneself — or to someone else. (Strawson 1966 : 165)

Strawson speaks of « ascribing a current or remembered state of consciousness to oneself ». This is potentially misleading, I think, for when one remembers, one is in a state of consciousness (a memory state) distinct from the state of consciousness one was in when one perceived the scene one presently remembers. So there are two types of self-ascription at stake in memory : one ascribes to oneself the earlier perceptual state, but one also ascribes to oneself the current memory state. Nor is this distinction limited to the case of memory : we find it also in the case of perception.

In perception, one attributes to oneself the conscious state one is in ; but there is another form of implicit self-ascription, corresponding to that which we have

discussed so far. The perceiving subject does not merely ascribe to himself or herself the property of being in a certain inner state when she, for example, sees a flower in front of her. She also ascribes to herself the world-involving property of standing in a certain relation to the flower — a relation which makes it possible for the flower to cause the current conscious experience which the subject implicitly self-ascribes. The two types of implicit self-ascription are hard to distinguish in the perception case, because the inner state we are aware of reveals the world to us and our situation in it. The state is, as it were, 'transparent', and our awareness of it is also, simultaneously, awareness of what it reveals to us. In the case of memory the distinction is made easier due to the temporal distance between the inner memory state and what it reveals. What the inner state reveals is how the world was and how we experienced it in the past.

Two grades of immunity

Both types of implicit self-ascription are immune to error through misidentification (Table 1). Consider memory first. If I remember something, I cannot be mistaken as to *who* is remembering, that is, as to who is undergoing the current memory experience. Nor can I be mistaken as to who experienced the remembered scene in the past. As Strawson says,

> it would make no sense to think or say : I distinctively remember *that* inner experience occurring, but did it occur to me ? (Strawson 1966 : 165).

So, corresponding to the distinction between the current memory experience, and the past perceptual experience which is remembered, there is a distinction between two forms of immunity to error through misidentification. *Grade-1 immunity* characterizes the subject's implicit self-ascription of the conscious state she is in, while *grade-2 immunity* characterizes her implicit self-ascription of the relevant relation to that which her conscious state represents. In the case of memory, the relevant relation to the scene remembered is the relation of *having experienced* that scene. So grade-1

immunity characterizes the subject's self-ascription of the memory state she is currently enjoying, while grade-2 immunity characterizes the subject's self -ascription of the property of having experienced the scene she remembers.

The subject's relation to the remembered scene, viz. the relation of having experienced it, corresponds, in the memory domain, to the subject's relation to the scene perceived in the perception domain. Just as the subject who remembers a flower implicitly ascribes to herself the property of *having seen* the flower she remembers, the subject who sees a flower implicitly ascribes to herself the property of standing in a certain causal-epistemic relation to the flower in front of her (the relation of *seeing* the flower). In the two cases the self-ascription is immune to error through misidentification (grade-2 immunity). The remembering subject cannot be mistaken as to who experienced the remembered scene in the past ; nor can the perceiving subject be mistaken when she ascribes to *herself* the property of standing in a certain relation to the flower in front of her. To be sure, she may not stand in that relation to anything if, for example, she is hallucinating and there is no flower there. But it cannot happen that she mistaken because she has misidentified as herself the person she knows to see a flower.

As I said, grade-1 immunity characterizes the other kind of implicit self-ascription. Both in the perception case and in the memory case, the subject self-ascribes the property of being in a certain conscious state : a perceptual state and a memory state respectively. The subject can no more be mistaken as to who is remembering, than she can be mistaken as to who experienced the remembered scene in the past. Similarly, the perceiving subject can no more be mistaken when she ascribes the current conscious state to herself than she can be mistaken when she ascribes to herself the relevant relation to the perceived flower.

	Perception	*memory*
Grade-1 immunity	The subject self-ascribes the perceptual state s/he is in	The subject self-ascribes the memory state s/he is in
Grade-2 immunity	The subject self-ascribes a certain relation to the	The subject self-ascribes a certain relation to the

	object of perception	object of memory

Table 1

       I do not wish to imply that the two types of self-ascription are separable ; for they are not. The self-ascription of the current perceptual/mnesic state and the self-ascription of the relevant relation to the perceived/remembered scene are two sides of the same coin. Yet they are conceptually distinct. As we shall see in Chapter 21, we can imagine fantastic situations in which one is legitimate while the other is not. In the light of such situations, one may argue that grade-1 immunity is stronger than grade-2 immunity. For the situations in question are situations in which the type of self-ascription which is grade-2 immune turns out to be false. Cases of quasi-memory and quasi-perception, to be discussed in the next chapter, are like that. The conceivability of such cases shows that grade-2 immunity is contingent or 'de facto'. It is much harder, if not downright impossible, to conceive of cases which falsify self-ascriptions displaying grade-1 immunity.

Chapter 22

Quasi-perception and quasi-memory

Wrong connections

I said that it is a contingent fact that information gained in the internal mode can only be about our own body. Indeed, we can certainly imagine a fantastic situation in which we would be connected to someone else's body and feel the condition of that body.[46] Similarly, it is a contingent fact that what we perceive is the world around us. We can imagine a fantastic situation in which we would perceive the world around somebody else. Evans describes a not-so-fantastic situation in which one hears sounds that are actually occurring in a remote place :

> It is possible that the subject is wearing a pair of ultra-lightweight and undetectable earphones, operating in such a way that the sounds he hears are not sounds at the place he occupies at all — they may be sounds occurring in some other place. (Evans 1982 : 184)

It is easy to imagine an analogous situation in which (perhaps because of undetectable glasses) what we see is not what is before our eyes, but what is before someone else's eyes.

---

[46] « The fact that [certain] methods of discovering a person's states are for the exclusive use of the person in question, is a matter of quite reliable but not quite necessary facts. We could imagine, for example, cases involving spinal columns that are connected across bodies which had the result that one knew about the state of another person's stomach in the way that we normally know only of the states of our own stomach. » (Perry 2000 : 335-336)

All such cases are cases of 'quasi-perception'. In quasi-perception, the subject is in a certain state, characterized by a content and a mode. The mode normally justifies the subject in implicitly self-ascribing a certain relation to the worldly condition that is the content of the state ; but what distinguishes quasi-perception from perception is the fact that the self-ascription may turn out to be *false* : it may be that, because a wrong connection is in place, it is not the subject, but someone else, who stands in the right relation to the worldly condition that is the content of the state.

The label 'quasi-perception' is formed after the celebrated expression 'quasi-memory' coined by Shoemaker to refer to an analogous predicament in the memory domain. In quasi-memory, as in memory, the subject is in a state M(p), where 'M' is the memory mode and 'p' is the propositional content of the state. In the actual world, a subject in such a state is thereby entitled to believe that he had a past perception that p, which perception left the memory as a trace. But in a world in which humans have quasi-memory, it may be that a wrong connection is in place : the subject who remembers may *not* be the subject who enjoyed the perception in the past. The memory is a trace of an earlier perception, yet the perception in question may have been experienced not by the subject himself, but by some other person.

What happens in quasi-perception and quasi-memory can be described by saying that there is *a failure of grade-2 immunity*. Grade-1 immunity is retained : the subject is entitled to automatically self-ascribe the property of being in a certain perceptual/mnesic state, without any possibility of error. What the subject cannot rightly do, however, is automatically self-ascribe the property of standing in the right relation to the content of the state : for it may be that he does *not* possess the property of seeing a flower in front of him, or of having seen the remembered scene in the past. Someone else may be seeing, or someone else may have seen. The subject entertains the state, yet what the state reveals need not concern the subject, but may concern someone else.

Failure of immunity ?

The situations I have described do not happen, but they *could* happen ; by which I mean that they are conceivable. Now that very possibility raises a problem. What is

to prevent us from construing perception as a particular form of quasi-perception, and memory as a particular form of quasi-memory ? In quasi-memory, the rememberer *may or may not be* the original experiencer. What Shoemaker calls the 'previous awareness condition' may or may not be satisfied. Now, in a normal case of memory, the rememberer *is* the original experiencer, but we have just conceded that *things might go otherwise* : the rememberer might not be the experiencer. Is this not sufficient to show that memory *is* a special case of quasi-memory, namely the case in which the original experiencer turns out to be the rememberer (but might not be) ?

If, following Shoemaker (1970), we accepted the conclusion that ordinary memory is (a special case of) quasi-memory, we would be faced with a serious problem. We would have to give up the claim that memory is immune to error through misidentification. As Shoemaker writes,

> The immunity of first person memory claims to error through misidentification exists only because remembering requires the satisfaction of the previous awareness condition … This feature disappears once we imagine this requirement dropped. Quasi-memory, unlike memory, does not preserve immunity to error through misidentification. (Shoemaker 1970/2003 : 28)

But if immunity disappears, then it can no longer be maintained that in perception and memory, the subject's self-ascription of the relevant relation to what his experience represents is identification-free. We will have to include an identification component in the content of the representation, thereby contradicting our claim that what is conveyed by the experiential mode need not be explicitly represented in the content.

There are more than one way of resisting this argument, however. First, as Evans pointed out, an error may arise (as in the wrong connection cases) without the erroneous judgment resting on an identification. I will argue this point below. Second and more important, we must understand the modality correctly. Cases of wrong connection are 'possible' only in the sense that, in some possible world *w*, they would happen. We should clearly distinguish that possible world *w*, in which wrong connections happen, from the actual world, in which they do *not* happen and memory

123

and perception *are* immune to error through misidentification. A case of 'ordinary memory', that is, a case in which the rememberer is the original experiencer, would be a special case of quasi-memory in *w*, but in our world it is not. One of the things that show this is precisely the fact that memory is immune to error through misidentification, while quasi-memory is not.

Quasi-perception in the actual world

One can object that the fantastic situations we talked about may well obtain, some day, in the actual world. Let us assume that this is right, and consider what would happen if, in the actual world, an ordinary subject found himself in such a situation.

Consider the first case of quasi-perception I mentioned, that in which the subject is connected to someone else's body, whose condition he can feel 'from inside'. Let us suppose that this actually happens to some actual person. The internal mode of representation normally justifies that person in self-ascribing the represented property (*having one's legs crossed*, say); but due to malfunctioning caused by the wrong connection, the self-ascription, though justified in view of the general design of the system, turns out to be false in that particular case. As Evans points out, a subject in this condition has the *illusion* that he stands in the right relation to that which his experience is about. In the case at hand, he has the illusion of being the person whose legs are crossed. He feels exactly as if his legs are crossed. And the subject's illusory judgment that his legs are crossed is not identification-dependent. The falsity of the judgment does not show that its content involves a mistaken identification component ('I am the person whose legs are crossed'), but simply that the subject's justification in making the identification-free judgement depends upon an assumption of normal functioning of the system — an assumption which may well be defeated (Evans 1982 : 187).

Arguably, the situation is different if the subject, A, *knows* he is connected to B's body. Knowing that the system does not work properly, he will fight against the illusion of being *F* when he internally feels *F*-ness. What can we say about such a case ? Should we say that, in that situation, A judges that *F*-ness is instantiated, and has to explicitly identify the person — B — who is felt (from inside) to be *F* ?

That is far from obvious. In a first stage, as we have seen, the subject will be under the illusion that *he* is *F*, and the illusion will persist even though the subject knows that it is an illusion. After some time, arguably, the illusion will weaken. Assuming the subject can get rid of it, there are several possibilities. The internal experience of *F*-ness may simply become meaningless for the subject, because it provides him with no information he can reliably use. Or A may be able to give it a new meaning, by exploiting the connection to B's body. For example, let us suppose that A has devices for remotely controlling B's body (as in a video game), and learns to react to B's felt *F*-ness by making B behave appropriately. The subject will progressively gain the ability to *displace himself*, as it were, into B's body : he will feel and act through that body.[47] This is easier to imagine if we suppose that the subject's perceptions are *all* quasi-perceptions, where it is B, not A, who is standing in the right relations to what is perceived. I would say that, in such a case, *the subject will continue to self-ascribe the properties represented in the internal mode*, but he will be displaced into B's body. The subject will adjust to his new situation, just as a subject with inverted glasses does, after some time.

Introducing the identification component

There is a complication if we suppose that A has merely acquired an 'auxiliary' self which coexists with the subject's main self, as a videogame character coexists with the player of the game. This version requires on a part of the subject an ability to switch from one perspective to the other. In such a case also, I would think, the subject will continue to 'self-ascribe' the relevant relation to what the experience represents, but there will be two different 'selves' for the subject to ascribe the relation to.

This is where an identification component comes into the picture. In the imagined situation, the subject has to *infer* which 'self' is at issue. This type of inference is familiar : we have something similar in memory. We know that a

---

[47] See Dennett 1978 : 312-315 for an elaborate scenario in which such displacements would take place.

remembered experience belongs to our past, but often we don't know *when*, in the past, the experience took place — we have to infer this. (This may not be the case when the memory is very fresh, and bears its freshness on its sleeves ; see footnote 48 below.) The same indeterminacy might affect a self-ascription : the subject might have to infer which of his two 'selves' stands in the right relation to the content of the current experience.

      Shoemaker's quasi-memory examples are analogous but more radical than that. We are to imagine a world in which having a memory means that *someone* has experienced what we remember, but the memory mode does not tell us anything regarding the person who had the original experience. That must be inferred from the content of the experience, just as we infer the specific time and place of the original experience. Evans argues that if we had only such quasi-memories, we could never evolve a concept of ourselves (Evans 1982 : 243). That may be right. But in any case, the memory mode would, in such a possible world, have a different semantics from that of the memory mode in the actual world. In a memory state M(p), the mode M means that the subject of the state has perceived that p in the past. In the counterfactual situation imagined by Shoemaker, the memory mode means that *someone* has had the corresponding perception in the past (and, possibly, that this perception has caused the present memory). This sheds no more light on the actual semantics of memory states than the fact that we can imagine a counterfactual situation in which the memory mode means not only that the rememberer has had a corresponding perception in the past, but also that he's had such an experience at a particular time and place, somehow encoded in the mode.[48]

---

[48] The 'freshness' of a memory is a qualitative feature that tells us something about the time of the original experience. This feature — arguably a parametric property of the mode — could be exploited in a more systematic manner than it is.

Chapter 22

Reflexive states

Perception and memory

In perception and memory, the subject of the state — the person to whose mental life the conscious state belongs — is identical to the person who stands in some relation R, dictated by the state (and in particular by its mode), to what the state represents. In visual perception, the relation R_{vis} is the causal-epistemic relation of *seeing* the state of affairs represented. In the internal (proprioceptive/kinaesthetic) mode, the relation R_{int} is the relation of *instantiating* the property that is represented, or, better perhaps, the relation of *being* the person whose property is represented. In episodic memory, the relation R_m is the relation of *having perceived* what the memory state represents.

The mode fixes not only the nature of the relevant relation R, but also the higher-order relation (viz. identity) between the person bearing that relation R to what is represented and the person undergoing the conscious state. Shoemaker's thought experiments involving quasi-memory, and the analogous thought experiments involving quasi-perception, are useful because they enable us to single out this aspect of the contribution of the mode, namely, the relation of identity between the subject of the state and what, for want of a better term, I will call the *R-filler* : the person who stands in relation R to what is represented. The subject of the state and the *R*-filler are not easy to tease apart, precisely because their identity is guaranteed by the architecture of the system and the nature of the state. Teasing them apart is what the thought-experiments do for us.

Because of the identity between the subject and the *R*-filler, perception and memory are *reflexive* states :[49] in perception, a subject *x* is aware that *he himself* is perceiving some state of affairs ; through memory, the subject is aware that *she herself* has perceived the state of affairs she remembers ; etc. In all such cases a

---

[49] The expression 'reflexive state' comes from Higginbotham (1995).

constraint which I will later refer to as the 'Reflexive Constraint' is satisfied, and it is satisfied in virtue of the mode of representation.

Intentions and emotions

When one intends to do something, one is in a state with a content (what one intends to do) and a mode (the intention mode). The content of the state is an action, or rather an action schema : one's doing a certain thing. Because that action is represented in the intention mode, the state will count as satisfied only if *the subject of the state* fills the agent role in the schema – only if *he himself* does the thing in question. As Shoemaker writes, « intentions have as their 'intentional objects' actions to be done by the very person who has the intention » (Shoemaker 1970/2003 : 47).[50] In other words, intention is a reflexive state.

Among the reflexive states are many emotional states, such as guilt, shame, or pride (Shoemaker 1970/2003 : 47 ; Higginbotham 1995 : 237). Here is Shoemaker again :

> The appropriate objects of remorse, and of a central sort of pride, are past actions done by the very person who is remorseful or proud, and the appropriate objects of fear and dread, and of delighted anticipation, are events which the subject of these emotions envisages as happening to himself. (Shoemaker 1970/2003 : 47)

---

[50] Besides this condition, which is common to the intention to do *x* and the desire to do *x* and makes them both reflexive states, there is an extra condition which is specific to intention (and is similar to what we talked about in connection with promises in chapter 16) : the subject of the state must do the action as a causal result of entertaining the state. It follows that « a positive intention is… always in part about itself. It is always the intention that it itself will lead in such and such a way to such and such a result » (Harman 1986 : 85 ; see also Searle 1983, ch. 3). So intention is, as it were, a doubly reflexive state. It is reflexive because it requires the represented action to be performed *by the subject of the state*, but also because it requires the action to be performed *as a causal consequence of the state*.

The words 'delighted anticipation' provide a ready-made analysis of the complex state they denote in two components : the subject's anticipation, and the emotion of delight which that anticipation causes. The component state of anticipation is already reflexive because « one can only anticipate what is conceived to lie in one's own future » (Higginbotham 1995 : 237). To be sure, we can imagine a state of 'quasi-anticipation' where the anticipating subject would be anticipating situations lying in someone else's future. Some fortune-tellers pretend to have this gift of quasi-anticipation. But the emotion of delight is conditional upon the identity between the subject of the complex emotional state and the *R*-filler. If some very pleasant situation lies ahead in my neighbour's future, I will not be as happy as I would if it lay in my own future :

> Even the most unselfish man, who is willing to suffer [so] that others may prosper, does not and cannot regard the pleasures and pains that are in prospect for him in the same light as he regards those that are in prospect for others. He may submit to torture, but he would hardly be human if he could regularly view his own future sufferings with the same detachment (which is not indifference) as he views the future sufferings of others. (Shoemaker 1970/2003 : 48)

The analysis of 'delighted anticipation' in two components applies to the other emotions discussed by Shoemaker and Higginbotham. Remorse, pride and guilt are emotions caused by memories of things done, where the memory mode guarantees that the remembered actions are actions *we* performed. If those actions did not belong to our own past but to someone else's past — if the underlying memories were not genuine memories but quasi-memories — the emotions in question, to the extent that they would arise, would presumably be weaker than they actually are.

Another emotional state to which the two-component analysis applies is relief. Like the state of delighted anticipation, *being relieved that p* is a complex state, comprising an emotion (relief) and a doxastic state (the belief that *p*) which causes the emotion (Higginbotham 1995). In Prior's example, 'Thank Goodness that's over',[51] what causes the emotion of relief is the belief expressed by 'that's over' : the

---

[51] See the quotation from Prior's eponymous article in Chapter 13, p. 00.

belief that a certain painful episode has reached its conclusion and (therefore) belongs to the past. The episode in question must have been painful to the subject of the relief state, and it must be past with respect to that very state. This shows that relief is another reflexive state.

<div align="center">Reflexivism</div>

Reflexive states, as I have characterized them so far, are reflexive in virtue of their mode : the mode secures the identity between the subject of the state and the *R*-filler. There is immunity to error through misidentification precisely because the subject's involvement is secured by the mode without resting on an identification.

There is an alternative analysis of reflexive states, however. According to that alternative analysis, a reflexive state is *a state whose content refers to the state itself*. Such an analysis has been put forward by Higginbotham in connection with Prior's 'thank Goodness' example.

The object of relief is essentially tensed, according to Prior. Knowing that the conclusion of the painful episode takes place at a certain time brings relief only if the time in question is thought of in a certain way, corresponding to the use of the present tense. Higginbotham fleshes out Prior's scenario by imagining someone's relief that his root canal is over (Higginbotham 1995 : 226-49). When leaving the dentist's chair, at time t^* (say, Friday at four o'clock), the subject thinks : 'Thank Goodness that's over'. What brings relief is not the thought that the root canal is concluded at t^* ; for the subject might have known previously that the root canal would be over by 4 o'clock, and that would not have brought the same sort of relief as the thought that the root canal *is* over. Only the tensed thought brings relief, because the tensed thought presents the time at which it holds as *the present time*.

In 'Tensed Thoughts' and subsequent papers, Higginbotham brings together the extensional analysis of tense discussed in Chapter 5 and Prior's view that the content of relief must be tensed. According to (a simple version of) the extensional account, the proposition expressed by a present tense sentence such as 'That's over' involves a temporal relation (simultaneity or overlap) between the time of the painful

episode and the time of utterance.[52] If the subject's relief is felt but not linguistically expressed, we must replace the time of utterance by the time at which the state of relief occurs. On Higginbotham's analysis, the subject in Prior's example is in a complex state of relief *s*, consisting of an emotion (relief) and a belief that causes the emotion, where the content of the belief is the proposition *that the painful episode s' is before s.* The state is reflexive (self-referential) simply because its content involves a reference to the state itself.

In 'Remembering, Imagining, and the First Person' (2003a) Higginbotham generalizes this analysis to *all* reflexive states. They are all states the content of which is a reflexive proposition involving the state itself as a constituent. Higginbotham's view is similar to that held by John Perry since the early nineties. According to Perry, essentially indexical beliefs — the beliefs which are the most closely tied to perceptual intake and behaviour — are beliefs the content of which is a reflexive proposition.[53] And of course, both Higginbotham's and Perry's views are similar to Searle's self-referential analysis of the content of experience. There is a school of thought here, which we may refer to as 'Reflexivism'. Among reflexivists, I should also mention the advocates of a generalization of Reichenbach's token-reflexive analysis, such as Manuel Garcia-Carpintero.

There is a sense in which the reflexivists are right : the complete content of a reflexive state — its truth-conditions — cannot be spelled out without reflexively referring to the state itself. But I object to the temptation to think that, in the narrow sense also, the content is self-referential.

In Part V, I already criticized Searle's reflexivist analysis of perception and memory : the self-referential element Searle rightly detects comes from the mode, I

---

[52] Higginbotham actually uses Davidsonian event variables : rather than directly establishing relations between times, the tenses establish temporal relations between e.g. the situation talked about and the situation of utterance.

[53] Perry distinguishes different levels of content. The reflexive level is one, to be distinguished from what he sometimes calls the 'official content'. In *Knowledge, Possibility and Consciousness*, Perry introduces a contrast between « believing a [reflexive] proposition P and having a belief a reflexive content of which is P » (Perry 2001 : 132). I will not discuss the subtleties of Perry's reflexivism here nor compare it in detail with Higginbotham's, which will be my main target.

argued, and does not affect the content of the state (in the narrow sense of 'content'). The same thing holds for relief. What 'That's over' expresses is a temporal proposition, true at any time t iff the painful episode is over at t. What brings relief is the subject's belief, at t^*, that this proposition holds. The subject's belief is indeed true iff the temporal proposition in question is true at the time of the thought episode, that is, at t^*; but this does not make t^* (or the thought episode) a constituent of the proposition that is the (narrow) content of the relief-causing belief. t^* only comes into the picture through the act of *asserting* that proposition, linguistically or mentally.

In Part VII I will expose the limitations of the reflexivist approach. This will enable me to emphasize the merits of my own approach, based on the SMR framework and the mode/content distinction. One of those merits, as I hope I have shown, is that it provides a straightforward explanation of the phenomenon of immunity to error through misidentification.

Part Seven

Relativization and Reflexivity

Chapter 23
The (alleged) reflexivity of *de se* thoughts

De re and *de se*

If I think of a certain object, *a*, that it is *F*, I have a *de re* thought regarding *a*. The thought is true iff *a* is *F*, regardless of how the object *a* is thought about. In contrast, a descriptive (or '*de dicto*') thought about *a* is such that the mode of presentation plays a truth-conditional role (in the terminology of Recanati 1993 : the mode of presentation is 'truth-conditionally relevant' in descriptive thoughts, and 'truth-conditionally irrelevant' in *de re* thoughts.) Even if *a* turns out to be the strongest man in the world, the (descriptive) thought that the strongest man in the world can lift 150 kilos has general rather than singular truth-conditions, hence it is not specifically about *a*. Instead of being true iff *a* can lift 150 kilos, it is true iff there is a man *x* such that for every man *y* distinct from *x*, *x* is stronger than (or, perhaps, at least as strong as) *y* and *x* can lift 150 kilos.

De se thoughts are a particular case of *de re* thought.[54] A *de se* thought is a *de re* thought about oneself, that involves a particular mode of presentation, namely a first person mode of presentation. (One may entertain a *de re* thought about oneself that does not involve such a mode of presentation, as in Kaplan's famous mirror example.[55]) As Frege wrote in 'The thought', « every one is presented to himself in a particular and primitive way, in which he is presented to no one else » (Frege 1918-19: 25-6). I call the 'special and primitive' mode of presentation which occurs in first person thoughts '**EGO**' or rather '**EGO**$_x$' where 'x' stands for the name of

---

[54] This might be disputed, in light of the distinction, to be introduced in the next chapter, between basic and nonbasic *de se* thoughts. Only nonbasic *de se* thoughts are a variety of *de re* thought, it may be argued. See Chapter 36 for a detailed discussion of Lewis's claim that (basic) *de se* thoughts are not a variety of *de re* thought, but the other way round.

[55] Kaplan's example involves a man pointing to himself in the mirror and saying (or thinking) 'His pants are on fire', without realizing that he is the man whose pants are on fire.

the person thinking the thought (for example '**EGO**$_{Lauben}$' in the case of first person thoughts about Lauben).[56]

Qua *de re* thought, a *de se* thought such as the thought that *a* would express by uttering 'I am F' is true iff *a* is F. Like all *de re* modes of presentation, the mode of presentation **EGO**$_a$ does not affect the truth-conditions of the thought it occurs in. If *a* points to himself in a mirror and says 'He is F', the *de re* thought he then expresses has the same truth-conditions as the *de se* thought *a* expresses by saying 'I am F': both thoughts are true iff *a* is *F*. But of course, the thoughts themselves are different, as they involve different modes of presentation — a first person mode of presentation in one case, a demonstrative, third person mode of presentation in the other. As a result, thinking one thought has very different behavioural consequences than thinking the other (see Castañeda 1999 and Perry 2000 for well-known examples).

Ascribing de re *thoughts*

When one *ascribes* a *de re* thought, the mode of presentation under which the ascribee thinks of the object (the *res* that his *de re* thought is about) may become truth-conditionally relevant. It becomes truth-conditionally relevant if the ascription is 'opaque', that is, if the ascriber implicitly specifies the mode of presentation under which the ascribee thinks of the object his thought is about. If I say 'Lois Lane believes that Superman can fly', I mean that she does so when she thinks of him *as* Superman, not when she thinks of him as Clark Kent. It is because the mode of presentation becomes truth-conditionally relevant in opaque ascriptions that one may simultaneously assert 'Lois Lane believes that Superman can fly' and 'Lois Lane does not believe that Clark Kent can fly'. (Those ascriptions would be inconsistent if

---

[56] This notation with subscripts, which enables us to distinguish the mode of presentation type (**EGO**) from its instances (e.g. **EGO** $_{Lauben}$) is Peacocke's. See Peacocke 1981, 1983. On the distinction between indexical modes of presentation considered as types, as instances, and as occurrences, see Recanati 2005b : 24-25.

they were interpreted as 'transparent', that is, as *not* specifying the mode of presentation under which Superman/Clark Kent is thought of.)[57]

I said that in an opaque ascription, the mode of presentation is *implicitly* specified. I said so because a sentence such as 'Lois Lane believes that Superman can fly' merely suggests, but does not explicitly say, that Lois Lane thinks of Superman as Superman. This is what makes the semantics of belief reports especially problematic (and interesting): on the one hand we want to say that the mode of presentation affects the truth-conditions ; on the other hand we want to say that it has the status of 'pragmatically imparted information' (Salmon 1986) and is not explicitly contributed by some lexical item or phrase or construction in virtue of the semantic rules of the language. Because of these features belief reports are among the phenomena that call for a revision of the standard views regarding the semantics/pragmatics interface (Recanati 1993, 2004).

Yet there are a few cases in which it seems that the mode of presentation is explicitly contributed in virtue of the semantic rules of the language. Ascriptions of *de se* thoughts are a case in point.

Ascribing de se *thoughts*

Let us consider the following sentences :

(1)   John$_i$ expects that he$_i$ will be elected
(2)   John expects that he himself will be elected
(3)   John expects to be elected

In (1), the ascription may be interpreted as transparent or opaque, depending on the context. Transparent reading : we are told that John expects, concerning some individual who happens to be John himself, that he will be elected, but we are not told how he thinks of that individual. Opaque reading : some particular mode of presentation is contextually specified, as in the Superman case. For example, it may

---

[57] I should be clear at this point that the *de re*/de dicto and transparent/opaque distinctions are orthogonal : there are transparent and opaque ascriptions of *de re* thoughts. See Recanati 2000a and 2000b, ch. 9.

be contextually clear that (1) is meant to report a *de se* thought of John's. What (1) contextually says, on that reading, is that John has an expectation he might express by saying : 'I will be elected'. But the contextually specified mode of presentation may not be a first person mode of presentation. It may be contextually clear that John's expectation concerns the man now speaking on TV, who (unbeknown to John) happens to be John himself. On that reading what (1) contextually says is that John expects *that man* (who happens to be himself) to be elected. Here John might express his expectation by saying 'That man will be elected', but not by saying 'I will be elected'. (We can construct a scenario in which John expects the man he sees on TV to be elected, but does not himself expect to be elected.)

It is often said, following suggestions made by Castañeda in the sixties (Castañeda 1999), that the emphatic reflexive 'he himself' in (2) disambiguates (1) and imposes the *de se* reading. According to Castañeda, or at least, according to some of his followers, (2) can only report a *de se* thought of John's, a thought John himself could express by saying 'I will be elected' (see e.g. Corazza 2004: 280-281). So the emphatic reflexive is a 'quasi-indicator' that not only imposes an opaque interpretation of the report it occurs in, but imposes a specific opaque interpretation such that the ascribee thinks of the object his thought is about in the first person way. On this view, (2) means that John expects John, thought of under the mode of presentation **EGO**$_{\text{John}}$, to be elected.

I do not think this is quite right. I agree with Castañeda and his followers that (2) *strongly suggests* that what is being reported is a *de se* thought, because it is hard to interpret the emphatic reflexive otherwise. But this interpretation is not, strictly speaking, mandatory. One can imagine contexts in which (2) would not report a *de se* thought, and that shows that 'he himself' is not a 'quasi-indicator' in virtue of the rules of the language, but simply has quasi-indicator *uses*, like the unadorned third person pronoun 'he'. If this is true — and I will not argue this point nor give examples here — then the difference between (1) and (2) is not as great as one may have thought.

Still, there is a way of disambiguating (1) and imposing the *de se* reading by purely linguistic means. As various authors have argued, following Chierchia (1989), what distinguishes sentence (3) from (1), and perhaps also from (2), is that (3) can *only* be interpreted as opaquely ascribing a *de se* thought.

Sentence (3) is usually analysed as involving a covert anaphoric element, PRO, as subject of the infinitive clause, with the subject of the main clause as antecedent. So (3) says that John expects PRO to be elected, where PRO inherits its reference from the antecedent 'John'. Besides being anaphoric on the subject of the higher clause, a second essential feature of PRO is that it imposes an interpretation of the complement clause as denoting a *de se* thought. Thus (3) reports a *de se* expectation, and (4) or (5), which use (almost) the same construction with PRO,[58] also ascribe to the subject, John, a state whose content is irreducibly first personal :

    (4)   John remembers PRO being elected

    (5)   John imagines PRO being elected

This can be seen by contrasting (4) and (5) with 'John remembers *himself* being elected' and 'John imagines *himself* being elected' respectively : these sentences *can* be given a non-*de se* reading, while (4) and (5) cannot (Higginbotham 2003a : 510-512). In (4) and (5), what is remembered or imagined can only be about *the subject*, i.e. the very person imagining or remembering, not only in the *de facto* sense that it is John whose election is remembered or imagined by John, but in the stronger sense that John, the person the memory/imagining is about, is thought of as *the person currently remembering or imagining*. This is where the *de se* element comes from, according to Higginbotham. It is not simply that John has a thought about John : he has a thought reflexively concerning *the thinker of that thought*. Not only is the individual the thought is about the same as the individual who has the thought, but that individual is *thought of as* the individual who has the thought. This suggests to Higginbotham the following analysis of PRO :

> Following the analogy of reflexive thoughts with respect to time, (…) we
> identify as the peculiar semantic contribution of PRO that it presents the

---

[58] I say 'almost' because the infinitival construction is replaced by a gerund. I will ignore this difference.

subject as the subject (or experiencer) of the event or state e as given in the higher clause, or $\sigma(e)$ for short. (Higginbotham 2003a : 514)

Higginbotham therefore analyses (3), (4) and (5) as (6), (7) and (8) respectively :

(6) $\exists e_1$ Expects [John, e_1, $^{\wedge}(\exists e_2)$ Is-elected $(\sigma(e_1), e_2)$]

(7) $\exists e_1$ Remembers [John, e_1, $^{\wedge}(\exists e_2)$ Is-elected $(\sigma(e_1), e_2)$]

(8) $\exists e_1$ Imagines [John, e_1, $^{\wedge}(\exists e_2)$ Is-elected $(\sigma(e_1), e_2)$]

Here 'e_1' and 'e_2' are Davidsonian event variables, which Higginbotham, like Parsons, uses to stand for states as well as for events properly speaking. Leaving temporal considerations aside, (6), (7) and (8) say that there is a state e_1 of expecting, remembering or imagining whose subject (experiencer) is John, and whose content is the proposition that there is a state or event e_2 of being elected whose subject/experiencer is *the subject of e_1*. Analysing the semantic contribution of PRO in this way guarantees both the anaphoric link between PRO and the subject of the higher clause and the first personal quality of the thought denoted by the complement clause.

<div align="center">De se thoughts as reflexive thoughts</div>

Higginbotham's analysis of ascriptions of *de se* thoughts yields an analysis of *de se* thoughts themselves. A *de se* thought is a thought containing the special mode of presentation **EGO**, which we can now analyse as a sort of mental PRO. To think of an individual as oneself (**EGO**), one has to think of him or her as the thinker of that very thought. Just as Prior's relieving thought 'That's over' is, for Higginbotham, a mental state e whose content is the reflexive proposition that the painful event — the root canal, say — is temporally located before e (Chapter 22), the *de se* thought expressed by 'I am F' is a mental event e whose content is the reflexive proposition that $\sigma(e)$ — the thinker of this very thought — is F.

Chapter 24

Reflexivity : internal or external ?

Two approaches

As I emphasized when I discussed Searle's views, I think whatever reflexivity characterizes experiential states such as perception or memory is external and comes from the mode of representation. The content of the representation, in the narrow sense in which 'content' contrasts with 'mode', is not reflexive. Thus I disagree with Higginbotham's analysis, which construes the relevant reflexivity as internal to content.

For Higginbotham, the subject who remembers having been elected entertains a memory e_1 the content of which is the following proposition :

(1) $(\exists e_2)$ Is-elected $(\sigma(e_1), e_2)$ & $e_2 < e_1$

The proposition in question is doubly reflexive : it refers to the subject who undergoes the memory state of which it is the content, and it locates the remembered event e_2 as anterior to that memory state ('<' is the relation of temporal precedence).

I would offer a much simpler analysis. I take the content of the subject's memory to be simply the event of being elected. In the Davidsonian framework favoured by Higginbotham, this simple content can be represented as (2) :

(2) $(\exists e)$ Is-elected (x, e)

This is both a temporal and a personal proposition.[59] (2) is the proposition *that one is elected* — a proposition which (given a fixed world) is true at all and only those

---

[59] A personal proposition is the counterpart of a temporal proposition in the personal domain : it is a proposition that is true at some persons, and false at others. The counterpart of an eternal proposition is an impersonal proposition. (The terminology personal/impersonal comes from Nagel 1970, chapter 11).

person-time pairs <*x, t*> such that *x* is elected at *t*. A memory with that proposition as content is true *tout court* iff the proposition is true with respect to the subject who remembers and a time anterior to the memory.[60] It follows that if John has, at time *t**, a memory whose content is (2), John's memory is true iff John has been elected prior to *t**. On this view, both the pastness of the remembered event and the subject's involvement in that event derive from the fact that what is represented is represented in the memory mode.

<center>*Basic and nonbasic* de se *thoughts*</center>

The same 'externalization' move can be made (and, I will argue, must be made) with respect to *de se* thoughts.

  In the SMR framework, there are two types of *de se* thought : basic and nonbasic. A nonbasic *de se* thought is a thought the (narrow) content of which involves an 'identification component', through which the object thought about is identified as oneself. The subject who sees himself in the mirror and thinks 'My legs are crossed' entertains such a thought, and explicitly thinks of the person whose legs are crossed as being herself. (As Evans would put it, the content of the thought is conjunctive : *Someone's legs are crossed, and I am that person.*) Here the concept EGO occurs on the side of the *lekton*, so the thought may be said to be *internally de se*. In contrast, basic *de se* thoughts are identification-free, and they are *de se* only externally: no concept EGO occurs as part of the *lekton*. The *lekton* is a personal proposition, without any constituent corresponding to the person to whom a property is ascribed. The first person comes into the picture only at the evaluation stage : to *assert* a personal proposition is to present it as true with respect to oneself and the present time. Since a personal proposition is a property of person-time pairs, for someone to assert such a proposition is, in effect, to *self-ascribe* the property that occurs on the side of the *lekton* (Loar 1976 : 358 ; Lewis 1979). Here, the self-ascription introduces an element of reflexivity (since the property that is the content of the assertion is ascribed to the person who makes the assertion) but it is external to content : it comes from the assertive mode.

---

[60] This is a simplification. One should add a causal condition : the remembered event must be the causal source of the current memory.

The importance of basic de se *thoughts*

The distinction between basic and nonbasic *de se* thoughts accounts for the observation that some, but not all of our *de se* thoughts have the property of immunity to error through misidentification. When, looking at the mirror and seeing my legs crossed, I think 'My legs are crossed', I entertain a *de se* thought — a thought about myself in which I think of myself 'in the first person' — yet the thought is *not* immune to error through misidentification. According to the SMR account, immunity to error through misidentification follows from the fact that (only) basic *de se* thoughts are identification-free, since they do not involve the concept of self on the *lekton* side.

The notion of a basic *de se* thought in which the self is not represented (Perry 1986b) is important not only to understand the phenomenon of immunity to error through misidentification, but also to understand the concept of self that occurs in nonbasic *de se* thoughts. Indeed, the ability to entertain basic *de se* thoughts is arguably a necessary condition for anyone to evolve the concept EGO. That is so because, as suggested by Evans, Perry, and myself following them, the concept EGO is best construed as a repository for information gained in a first person way (a repository which, in virtue of the Generality Constraint which applies to concepts, is also hospitable to information concerning the same object as information gained in the first person way — see Recanati 1993 : 123-25). Now a piece of information is gained in the first person way if and only if it is the content of a basic *de se* thought. It follows that the first step in an elucidation of the concept of self is a correct analysis of the functioning of basic *de se* thoughts.[61]

---

[61] Nozick criticizes the approach to *de se* thoughts in terms of an internal (first person) way of gaining information, on the grounds that (in my terminology) not all *de se* thoughts are basic. « We might imagine there is some way of observing ourselves which cannot be used to observe anything else. On this view, I know it is I who is in pain, for example, by observing in that particular way that someone is in pain. (…) This does not fit easily knowing nonpsychological statements that are reflexively self-referring (as my knowing I was born in Brooklyn, New York). (…) Reflexive access to ourselves, then, cannot be a special mode of relating to ourselves as objects » (Nozick 1981 : 81). I dispose of this objection by taking

As we saw in Chapter 23, Higginbotham treats the concept **EGO** as a mental equivalent of PRO analysable as 'the thinker of that very thought'. One advantage of this analysis is that it makes the mental indexical **EGO** similar to the linguistic indexical 'I', and extends to the former Reichenbach's token-reflexive analysis of the latter. Evans criticized this extension of the Reichenbachian analysis as illegitimate, but I will not press this line of argument here.[62] Rather, the objection I will raise concerns the distinction I have just introduced between basic and nonbasic *de se* thoughts.

That distinction is missing in the reflexivist framework as spelled out by Higginbotham : all *de se* thoughts are seen as involving the **EGO** concept, that is, σ(e). How, then, can the phenomenon of immunity to error through misidentification — and its failure in an entire class of cases — be accounted for ? Higginbotham cannot account for it by appealing to the identification-free character of basic *de se* thoughts, since he does not acknowledge the existence of such thoughts.

Higginbotham offers another explanation of the phenomenon. His discussion of immunity to error through misidentification is concise, and I am not sure I understand him correctly, but my (tentative) interpretation of what he says is the following : The concept 'the subject of this state' that the subject uses in thinking about himself is used, as it were, *attributively* when the thought is immune to error through misidentification: the subject refers to the subject of the state, whoever s/he is. Since that is so, no identification mistake can be made. In the famous Donnellan example, no error of identification can be made by saying 'Smith's murderer is insane' when 'Smith's murderer' is used attributively (Donnellan 1966). Errors of identification can occur only if one uses the description *referentially* to talk about, say, Jones : in such a case the speaker exploits the additional premiss that Jones is

---

basic *de se* thoughts as explainable in terms of the internal mode, and nonbasic *de se* thoughts as explainable in terms of basic *de se* thoughts.

[62] In a nutshell, Evans points out that the ability to demonstratively individuate a thought cannot be presupposed to be independent of the ability to entertain thoughts about persons in general, and *de se* thoughts in particular.

Smith's murderer, and that premiss may turn out to be false. In attributive uses there is no such additional premiss, so the judgment is, in a certain sense, identification-free even though, in a weaker sense, the object the thought is about is 'identified' as Smith's murderer.

On Higginbotham's account as I understand it, failures of immunity in first person thoughts are similarly traced to the fact that the subject does not primarily or exclusively think of the person the thought is about as σ(e), but in some other way. The use of the first person in such cases is justified by the additional belief that the object identified in that other way is identical to σ(e). So a *de se* belief is involved, and *it* exhibits the property of immunity which derives from employment of the reflexive concept σ(e), but it is only one element in a complex belief state which involves also non-*de se* elements (such as the belief that *the person in the mirror* has his legs crossed). If the speaker only thought of *himself*, i.e. σ(e), as having his legs crossed, his self-ascription would be immune to error through misidentification ; but the subject who sees himself in the mirror also thinks of himself as *the person in the mirror,* and that is the source of the failure : the judgment that the person in the mirror is σ(e) may well be mistaken, hence the *de se* belief based on it — the belief that σ(e)'s legs are crossed — is susceptible to error through misidentification. In attributive uses the *de se* judgment is not based upon such an additional premiss, so the judgment is, in a certain sense, identification-free even though, in a weaker sense, the object the thought is about is 'identified' as σ(e).

Chapter 25

What is wrong with Reflexivism

Higginbotham and the blindsighted neuroscientist

Let us consider the following passage where Higginbotham summarizes his views regarding *de se* thoughts and their immunity to error through misidentification:

> What is the reason for immunity to error through misidentification in the case of thinking, on the basis of a present perception, 'I hear trumpets' ? I shall assume it is this : that when I am in the relevant perceptual state, what I think is *that the subject of that state hears trumpets*. Hence, there can be no question of my *identifying* myself as the subject of the state. (Higginbotham 2003a : 507)

The example of the blindsighted neuroscientist from chapter 20 is a prima facie counterexample to this analysis. The blindsighted neuroscientist is, and knows he is, in a visual state e_1 which, because of his condition, is not available to his consciousness. Knowing that he is in that state, and thinking of himself as the person currently in that state, he conjectures (from the sounds he is hearing) that what he sees is a canary. So he has a thought the content of which is : that the subject of e_1 — the visual state he is actually in — is seeing a canary. It seems that Higginbotham's conditions are satisfied : the subject is in the relevant visual state, and he thinks the subject of that state sees a canary. Yet, as we have seen, the judgment is clearly not immune to misidentification if it is formed on the basis of the sounds the subject is hearing (at the same time as he blindly sees the canary).[63]

---

[63] Higginbotham might argue that a crucial condition is not satisfied in this case: the description 'the subject of e_1' is not used attributively, since the subject intends to refer to himself, and picks this mode of presentation only because he believes he is the subject of e_1. But that objection is not convincing. Since 'e_1' is characterized as whatever visual state

In response, Higginbotham might point out, correctly I think, that we do not have genuine reflexivity in that example, since the subject's thought that the subject of e_1 sees a canary is not the same mental state as e_1 itself. e_1 is the visual state of (blindly) seeing the canary. The thought is simultaneous with the visual state and accompanies it, but it is not identical to it.

In the ordinary cases Higginbotham has in mind — when, for example, the subject hears trumpets and judges that he hears trumpets — there is also a distinction between the perceptual state and the judgment. The judgment accompanies the perceptual state and is, perhaps, an aspect of that state, but it is not identical to it. Still, there is a relevant difference between a normal episode of hearing trumpets and the rather abnormal situation of the blindsighted neuroscientist who judges from the sounds he hears that he is seeing a canary. In the former case, it makes sense to say that the judgment that the subject hears trumpets is an aspect (a part) of his auditory experience. This is enough to render the experience reflexive since the judgment, which (on this assumption) is part of the overall experience, refers to the subject of the experience. In the blindsighted neuroscientist case, in contrast, the subject's judgement that he sees a canary cannot be considered as an aspect (a part) of his perceptual experience. The subject's judgment is clearly not part of his auditory experience, since it is inferentially based on it. Nor is it part of his visual experience — since, not being conscious of seeing, the subject does not have anything that may count as a 'visual experience'. So we may accept that in a normal perceptual episode, the judgment that one is perceiving what one is (conscious of) perceiving is an aspect of the perceptual experience, thereby guaranteeing reflexivity (the judgment refers to the experience of which it is a part). No reflexivity is achieved in the blindsighted neuroscientist case because the visual state is not conscious and the subject's judgment that he is seeing a canary is not an aspect of his visual experience. The judgment refers to the subject of the visual state, but the visual state is disjoint from the conscious mental state to which the judgment belongs.

(token) the subject actually is in, 'the subject of e_1' can only refer to the blindsighted neuroscientist ; it cannot refer to anybody else. So we can stipulate that what the blindsighted neuroscientist thinks is : 'The subject of e_1, whoever he is, is seeing a canary', since that will make no difference whatsoever. Here the description 'the subject of e_1' is used attributively, but immunity is still not achieved.

The crucial objection

Even though the counterexample ultimately fails, it reveals what I take to be a fundamental problem for Higginbotham's account (and Reflexivism more generally). According to the reflexivist account, to have a *de se* thought is to be in a mental state the content of which is the proposition that the subject of that state is *F* (for some value of '*F*'). Similarly, to have a tensed thought is to be in a mental state the content of which is the proposition that the time of the state is *F*. Mental indexicality thus reduces to reflexivity. To this attempted reduction I object that Higginbotham's conditions are not sufficient. A further condition must be satisfied, without which the account does not work : Not only must the state e_1 the subject is in have a reflexive content (including a reference to that very state), but it must be a *fully conscious state*. That means at least that (i) the state must be conscious (the subject must be experientially aware of it), (ii) the subject must be aware of it as a state of his own (the subject must self-ascribe the state), and (iii) he must be aware that he is presently experiencing the state. If any of these conditions fails, that is, if the state e_1 the subject is in is not fully conscious, then, even if its content is reflexive in the right way, the state will not count as a *de se* thought or a tensed thought or whatever it is that Higginbotham is trying to analyse.

If justified, this objection is rather devastating, for the extra condition required to make the analysis work also makes the analysis superfluous. In order to analyse what it is for a thought to be about oneself or about the present, Higginbotham presents complex conditions on the content of the mental state, but then it turns out that there is an additional condition : the state itself must be fully self-conscious, that is, the subject must think of it as a state *he himself* is in at the *present* time ! The entire *analysandum* reappears in the extra condition that must be added to the *analysans* to make it work.[64] No objection, I think, could be more devastating than this one.

---

[64] As Nozick puts it, « reflexive self-knowledge is not merely a person's knowing that the semantics of self-reference holds of some (thought or sentence) token. We need to add the very phenomenon to be understood : his knowing that he himself produced the token » (Nozick 1981 : 80-81).

Of course, I have only stated the objection. I have not provided any support for it. The support actually comes from the counterexamples which the objection inspires. The counterexamples in question are examples in which the subject is in a certain state with the sort of reflexive content which Higginbotham presents as sufficient to get the right type of thought, but where the state is not fully self-conscious. If, in such circumstances, we do not get the right type of thought, this is evidence that Higginbotham's account fails and that the extra condition required to make it work (self-consciousness) makes the whole account superfluous.

The blindsighted neuroscientist case was an attempt at constructing such a counterexample. It was a case in which the mental state the subject is in is not a conscious state. The attempt fails because, as we have seen, the proposition which is supposed to provide the reflexive content of the state is not really reflexive since the state it refers to is disjoint from the state to the content of which it belongs. So the counterexample fails, but the general objection stands, and we only have to construct other counterexamples : counterexamples in which the state is conscious, but not fully self-conscious in the sense required for Higginbotham's account to deliver the right results.

De se *thoughts and schizophrenia*

Let us suppose that the state e_1 the subject is in *is* conscious, in such a way that its content (to the effect that the subject of e_1 is F) is truly reflexive, contrary to what happens in the previous example ; but let us suppose the subject is a schizophrenic patient who believes that, among the mental states he is conscious of, some are not really his mental states, but those of some other person that have been somehow implanted in him. Let us refer to them as the subject's 'alien' states, and, for any such state, to its putative owner as 'the Other'. Suppose the subject takes the state e_1 to be an alien state. e_1 might be, for example, the following thought passing through the subject's mind : 'The owner of this mental state is good and omnipotent' (or equivalently : 'I am good and omnipotent'). The schizophrenic subject in whose mental life this thought occurs will understand that *the Other* — the person from whom the thought emanates — declares himself/herself to be good and omnipotent. If he thinks his psychiatrist is the Other, he will tentatively ascribe the property of being good and omnipotent to the psychiatrist, not to himself. This is a

counterexample, because the subject is in a certain state (he consciously entertains a certain thought), the content of the state reflexively refers to the subject of the state (whom the deluded subject takes to be different from what it actually is), yet the subject does not entertain a *de se* thought, to the effect that he himself is good and omnipotent. What this shows is that being in a state with a reflexive content is not sufficient to ground a self-ascription, let alone an immune self-ascription. The state, as we have seen, must be conscious, but even that is not sufficient : the subject must not only be conscious of the state, he must be conscious of *being the subject of the state* (the 'owner' of the thought).

Higginbotham briefly considers the type of case I have just discussed :

> Could a person x be in a state e of imagining being F without *recognizing* that $x = \sigma(e)$, the subject of the property of events being imagined ? If this can happen, then perhaps, as tentatively suggested in Campbell (1999), that person would have thoughts of which he did not seem to himself to be the author. In any case, it seems safe to assume that any such condition would be pathological. (Higginbotham 2003a : 520)

Pathological, indeed ; but that does not save Higginbotham from the counterexample. What the example definitely shows is that it is *not* sufficient for entertaining a *de se* thought to entertain a thought with a reflexive content involving the concept $\sigma(e)$. An additional condition must be satisfied : the subject must self-ascribe the property of being the owner of that reflexive thought. That self-ascription, which Higginbotham's account presupposes, is *external* to the reflexive content in terms of which Higginbotham attempts to account for *de se* thoughts. To be sure, that 'external' self-ascription is automatic for normal subjects, and fails only in pathological cases ; but its existence is sufficient to refute Higginbotham's analysis of *de se* thoughts in terms of a purely internal property of their content (its alleged reflexivity).

Richard's counterexample

In 'Objects of Relief', Mark Richard has constructed a similar counterexample to Higginbotham's analysis of tensed thoughts as reflexive thoughts referring to the time of the thought (Richard 2003b). Richard's scenario is the following:

Suppose that I am given to particularly vivid recall of things that have happened in the past. (…) Not only do I on occasion so recall external events, but on occasion I have, or at least believe that I have, such recollections of my internal reactions to them. For example, I at times recall (without trying) hearing the tolling of the midnight bells and occurrently thinking, 'This [the bell tolling] marks midnight'. On some such occasions, I take the entire episode to be a memory, of my experience of the bells and my mental reaction thereto ; on others, I am unsure whether I am recalling something past or not. (Richard 2003b : 168)

In this setting, imagine the following situation : Richard hears a sound with characteristics M, and as part of his auditory experience *e,* judges that there is a sound with characteristics M audible at the time and place of *e*. However, because of his special condition, Richard mistakes his current perceptual experience (including the reflexive judgment that is part of it) for a vivid memory of an earlier perception (including the reflexive judgment that *was* part of it). He wrongly thinks he is recalling an earlier perceptual experience rather than currently having one. In such a case, Richard will fail to think of the sound as present, even though, in having the auditory experience *e*, he thinks of the sound as simultaneous with *e* (which he mistakenly locates in the past.) This establishes that, to think of something as present, it is not sufficient to think of it as simultaneous with the state of one's so thinking ; for one may temporally mislocate the state of one's so thinking. As Richard concludes

What makes Higginbotham's account plausible is the fact that a reflexive state indeed is *normally* one that presents itself as present. But, as I hope the argument has made clear, thinking reflexively of a state is only normally, not necessarily, thinking of it as present. (Richard 2003b : 170)[65]

In his response to Richard, Higginbotham surprisingly concedes the point. He accepts that his conditions are not sufficient, and that the extra condition which must

[65] See also Mellor 1998 : 42, on the 'self-intimating' character of experience and its role in the analysis of tensed thoughts.

be added to them (self-consciousness) makes the whole analysis at least *look* superfluous :

> The feeling of relief presupposes an element of self-consciousness, an element that is masked when one considers only public utterances, which by their nature are self-conscious acts. (…) The state of being relieved involves the state itself as a constituent of the object of relief, *and demands also the location of that state as present*… But now that the latter point has been made explicit, it can be questioned whether the apparatus of cross-reference was essential to begin with. (Higginbotham 2003b : 193-4)

Higginbotham still defends his analysis by arguing that it provides a key to the linguistic phenomenon of 'sequence of tenses'. Be that as it may, I take Reflexivism to be refuted by the counterexamples. Other phenomena will be discussed in Part VIII (Chapters 26 and 27), which can be accounted for in the SMR framework but arguably not in a reflexivist framework.

Part Eight

The first person point of view

De se thoughts and subjectivity

Three types of case

We have distinguished having a *de re* belief about an individual *x* who happens to be oneself, and having a true *de se* belief. In Kaplan's example, the subject who (unbeknown to him) points to himself in the mirror and says 'His pants are on fire' has a *de re* belief which is about himself, as a matter of fact, but which is not a true *de se* belief. When he realizes that *he* is the person in the mirror, he stands corrected : '*My* pants are on fire', he now says, and that expresses a genuine *de se* thought. The *de se* thought in question is not basic, however : it rests on an explicit identification (the person in the mirror = **EGO**). Even if one does not accept the SMR framework with its distinction between basic and nonbasic *de se* thoughts, there is an empirical difference between two sorts of *de se* thought : those that are immune to error through misidentification and those that are not. However, I will assume the SMR framework and continue to refer to such thoughts as 'basic' and 'nonbasic' *de se* thoughts.

So we end up with a double distinction : (i) a first distinction between *de re* thoughts about oneself that are not *de se*,[66] and true *de se* thoughts ; and (ii) a second distinction between nonbasic *de se* thoughts, which are subject to error through misidentification, and basic *de se* thoughts, which are immune to such errors. (See Ludwig 2005 : 111-13 for a clear statement of the resulting three-fold distinction.)

De se reports and their readings

There is a clear semantic difference between (1) and (2) :

---

[66] I shall henceforth refer to them as '*de re* thoughts' tout court, when no misunderstanding is to be feared.

(1)   He remembers himself delivering a speech to the salesmen

(2)   He remembers delivering a speech to the salesmen

The difference can be summarized as follows : Reports using the PRO construction, like (2), can only be reports of basic *de se* thoughts, while reports using a reflexive pronoun, like (1), are 'unmarked' and can be interpreted in any of the three ways.

Suppose the subject, A, was filmed delivering a speech to an assembly of salesmen. He retained no memory of that experience, but, having recently seen the film, remembers the episode from the film. At this point there are two possibilities : he may have identified himself as the person in the film, or he may not. If he has, this is a 'nonbasic *de se*' type of case. The thought is *de se*, but subject to error through misidentification. (The subject may have been wrong in identifying himself with the character in the film.) If he has not identified himself as the person in the film, this is a *de re* type of case : the subject who remembers the scene in the film remembers himself giving the speech, without realizing it is himself that his memory (and the film) is about. Sentence (1) is compatible with both cases, and it is also compatible with the 'basic *de se*' type of case, that is, the case in which the subject *has* retained first person memories of the original experience, without the mediation of the film.

In contrast to (1), which has all these interpretations, (2) can only have a *de se* interpretation. The *de re* interpretation is ruled out. Moreover, it seems that only the basic *de se* interpretation is possible.[67] So the semantic distinction between the two constructions does not correlate in any simple way with either of the two distinctions

---

[67] Kendall Walton may be interpreted as denying this in the following passage : « When Gregory imagines playing in a major league baseball game and hitting a home run, he may imagine this from the inside… But suppose he imagines hitting the home run from the perspective of a spectator in the stands. He visualizes the scene from that point of view, and his image of the field includes Gregory as he slams the ball over the center field fence and rounds the bases. This imagining is, I believe, best classified as *de se*. It is perfectly natural to describe Gregory as imagining *hitting* a home run » (Walton 1990 : 30). For me, this is not a *basic de se* imagining, and I would be reluctant to use the PRO construction. I would rather say that Gregory imagines *himself* hitting a home run. (For more on *de se* imagination, see below, Chapters 27-29.)

we made : that beween *de re* and *de se* thoughts, and that beween two types of *de se* thought. The first construction is unmarked and compatible with *all* interpretations (including the *de re* interpretation) while the second construction specifically requires the basic *de se* reading.

First person and third person point of view

There is yet another distinction : between the first person or subjective point of view and the third person or objective point of view. Both the *de re* case and the nonbasic *de se* case involve the third person point of view. The subject who remembers himself from the film, whether or not he has identified the character in the film as himself, has a third person point of view on the scene : he sees himself objectively from outside, as a spectator does. In contrast, the subject who directly remembers giving the speech to the salesmen, without the mediation of the film, has first person memories. He subjectively remembers what it was like to give the speech. (See Figure 1)

Figure 1

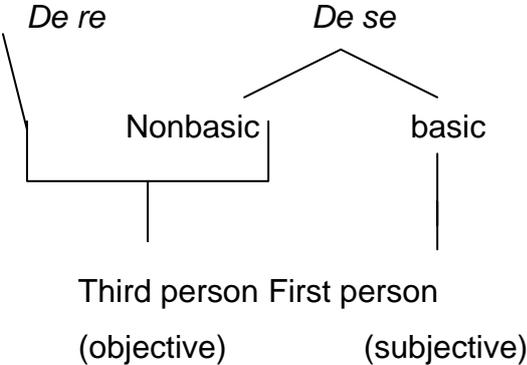According to Vendler (1979), both types of case are, in a certain sense, subjective. Episodic memory is *always* memory of an experience, hence something subjective and perspectival. But in one case the experience is, directly, the experience of, say, giving the speech to the salesmen, while in the other case (the 'objective' case) the experience is that of *seeing* the character in the film giving a

speech to the salesmen. Both experiences are intrinsically subjective, but in one case the character giving the speech stands as an 'object' in the subjectively experienced scene, while in the other case the person giving the speech is not 'in' the remembered scene : he is (only) the experiencer — the person whose experience is remembered — and has the same status as that of the spectator of the film.

Accounting for subjectivity

The PRO construction, which imposes the 'basic *de se*' reading, also imposes the first person point of view. Why? In Higginbotham's framework, the connection between on the one hand immunity (which distinguishes basic from nonbasic *de se* cases) and, on the other hand, first person point of view, is stipulated. Higginbotham assumes that the semantic contribution of PRO is a complex mode of presentation analysable in two distinct features : one feature (which Higginbotham represents as '$\sigma(e)$') accounts for immunity, while the other (which Higginbotham represents as '$\theta(e')$') accounts for the subjective or first person point of view :

> When one remembers or imagines *PRO being F*, then the linguistic element PRO is distinguished in two ways : (i) by being understood as the thing $\sigma(e)$ that is in the state *e* of remembering or imagining itself ; and (ii) by being at the same time understood as the bearer of the thematic role $\theta(e')$ as determined through the selection for the subject of the predicate *F(e')* (so in general that $\sigma(e)= \theta(e')$ is presupposed). I shall abbreviate this dual role of PRO as '$\sigma(e)$ & $\theta(e')$'. (…) We thus bring about the fact that… the subject cannot make an error of misidentification, and the fact that what is remembered [or imagined] is remembered [or imagined] *as* an action performed. (Higginbotham 2003a : 518)

This, however, is a purely descriptive move. No explanation is provided for there *being* such a conjunction of features. The conjunction is taken as a brute fact.

In the SMR framework an explanation is provided. Immunity and subjectivity are seen to derive from a common source. In the basic *de se* reading the *lekton* is a

personal proposition. This guarantees immunity (since no identification component is involved). But it also guarantees the subjective point of view : not being a constituent of the scene represented, the subject himself cannot be construed objectively and can only have the role of the spectator, i.e. that of the person from whose point of view the scene is experienced.

Chapter 27
Memory and the imagination

Crossmodal comparison

The typology of cases I have just presented in chapter 26 is appropriate for the memory mode. If we change the mode, the typology will change accordingly. So, for example, if we leave memory and turn to the imaginative mode, which is similar to memory in many respects, we find a significant difference.

Let us consider the similarity first. We find the same type of contrast between (3) and (4) as we found between (1) and (2), repeated below :

(1)   John remembered himself giving a speech to the salesmen

(2)   John remembered giving a speech to the salesmen

(3)   John imagined himself giving a speech to the salesmen

(4)   John imagined giving a speech to the salesmen

Like (2), (4) can only have a reading where what is imagined is imagined 'in the first person way'. The imaginer puts himself in the shoes of the person giving the speech and feels, or tries to (imaginatively) feel, like that person undergoing the relevant experience. First person imagination is just like first person memory.

In the case of memory we saw that, in contrast to (2), (1) was compatible with 'third person' memories mediated by the film. This seems to be true of (3) as well, in contrast to (4). I can certainly 'imagine myself giving the speech' from a third person, external point of view — by having the mental camera look at me from outside, as it were. This is phenomenologically different from imagining *giving* the speech.

At the beginning of his paper on imagination, Vendler emphasizes the phenomenological difference between third person and first person imagination, corresponding to the contrast between (3) and (4) :

We are looking down upon the ocean from a cliff. The water is rough and cold, yet there are some swimmers riding the waves. « Just imagine swimming in that water » says my friend, and I know what to do. « Brr ! » I say as I imagine the cold, the salty taste, the tug of the current, and so forth. Had he said « Just imagine yourself swimming in that water », I could comply in another way too : by picturing myself being tossed about, a scrawny body bobbing up and down in the foamy waste. In this case, I do not have to leave the cliff in imagination : I may see myself, if I so choose, from the very same perspective. Not so in the previous case : if I indeed imagine being in the water, then I may see the cliff above me, but not myself from it. » (Vendler 1979 : 161)

This illustrates what, following Vendler, I said earlier about the subjective and the objective : the objective imagination is a particular case of the subjective, namely the case in which the subject imagines *seeing* himself swim in the water. In that special case the subject plays two roles : he is not only the experiencer, the person from whose point of view the scene is seen, but he is also an object in the scene. This duality enables the subject to look at himself (herself) from an external, third person point of view. In contrast, the subjective case is the case in which the subject only plays the role of experiencer : s/he is not an object, an aspect of what is represented. What is represented is only « the cold, the salty taste, the tug of the current, and so forth ».

So much for the similarity between imagination (and imagination reports) and memory (and memory reports). Now, to the difference.

World-involvingness

What one's perception or memory is about is often determined by external factors of which we are unaware, like the identity of the object who stands in the right causal-epistemic relation to us. With imagination, the situation is different. What determines what one's imagination is about is not some external fact of which one may be unaware, but what Bernard Williams aptly called one's 'imaginative project' (Williams 1973). So, when we turn from perception and memory to imagination, we loose the *de re* case, that is, the case in which I represent something about myself without

realizing that it is myself whom the representation is about. It follows that (3) does not have as many interpretations as (1). I can hardly 'imagine myself being F' without realizing that it is myself whom I am imagining being F.[68]

Can we, in imagination, get the counterpart of the nonbasic *de se* case ? I think we can, but again, there will be a difference. It will not be a matter of the subject's wrongly or rightly identifying as himself the person his imagination is about since, as we have just seen, what one's imagination is about is a matter of intentional stipulation and does not have to be identified.[69] To get the counterpart of the nonbasic *de se* case, rather, we must exploit the distinction between the first person and the third person point of view which goes together with that between basic and nonbasic *de se* cases.

In basic *de se* memories, as we have seen, what one remembers is viewed from inside : we remember the action as it appears to the agent who performs it. In *de re* and nonbasic *de se* memories, what one remembers is viewed from outside :

---

[68] If I imagine the winner being F, and if, unbeknown to me, I am the winner, do I imagine myself being F ? I do not think so. A more problematic type of case occurs when the thinker's imaginative project itself is *de re*. Thus Walton writes : « Wittgenstein observed that when a person imagines King's College on fire, there may be no room for questioning his claim that it is King's College which he imagines — even if another college or a Hollywood movie set perfectly matches his visual image (if he has one). Doubt about the identity of the imagined college is not entirely out of the question, however. If the imaginer has prevously mistaken another college for King's, and if his intention is to imagine *that* college, he may think the college he is imagining is King's when it is actually the other one. If he saw King's College without realizing it was King's, he might then imagine King's College on fire without realizing that King's is the college he is imagining » (Walton 1990 : 29-30). Similarly, as John Hawthorne pointed out to me, a subject may entertain a *de re* thought about himself (as in Kaplan's mirror example), and imagine *that person* being F : that makes his self-imagination *de re*, though in a derivative manner. The claim that imagination does not allow for the *de re* case should therefore be qualified : no *de re* case is possible *unless* the *de re* character of the imagining is inherited from a mental state in another mode than the imagination mode, through some kind of anaphoric link. In Walton's example, the imagining is anaphoric on an earlier visual state.

[69] That is so even in daydreaming, an imaginative activity in which the will plays no part.

we remember someone doing the action and (in the nonbasic *de se* case) we identify that person as ourselves. The same thing is possible with imagination : as Vendler emphasizes in the quoted passage, we may imagine something about ourselves by adopting an external point of view — the point of view of an outside observer. This comes as close to the nonbasic *de se* case as is possible in the imaginative mode.

Conclusion

We have found that there are only two possible interpretations for a statement like (3) : we may imagine ourselves giving the speech, seen from an outside point of view, or we can imagine giving the speech, from the internal point of view. This distinguishes (3) from (1), which has an additional *de re* reading.

This distinction is entirely due to the difference between the memory mode and the imaginative mode. Like the perception mode, the memory mode is world-involving in a way in which the imaginative mode is not ; so it is possible to be mistaken as to what one's perception or memory is about, but no such mistake is possible in the case of imagination, which is a purely intentional activity.

Despite what I have just said, we cannot simply conclude that imagination reports have fewer interpretations than memory reports. It is true that the *de re* reading is missing in the imagination case, but the difference between the imaginative mode and the basic experiential modes (perception and memory) also makes an extra interpretation possible in the imagination case, which is not available in the other modes. That extra possibility, characteristic of imagination, shows up when we consider (4) : I will argue that, in contrast to (2), (4) can be interpreted in two distinct ways. One interpretation, which I call the 'quasi-*de se*', has no counterpart in the case of (2).

Chapter 28
Imagination and the Self

Being Napoléon

Can we imagine being Napoléon ? Of course we can. Yet this raises a problem, if one admits, as a piece of metaphysical truth, that one could not *be* (nor have been) Napoléon. There are accidental properties which one has and which one might not have, but one's identity is not among them. I am François Recanati, Napoléon is another person, living in a different century, and I could simply not have been him, or be him — although I might have had different properties than those I actually have. I might have been a high-ranking general, or an Emperor. But I might not have been Napoléon himself. Still, I can easily *imagine* being Napoléon.

Vendler says the tension between metaphysical impossibility and imaginability is relieved if we pay attention to the difference between imagining *oneself being Napoléon*, and imagining *being Napoléon*. Only in the former case is something impossible imagined :

> I can imagine being you, Napoléon, or even Napoléon at the battle of Waterloo. (…) « But », you object, « how can you imagine being me or, worse, being Napoléon, since you are not and this, as we are told, is a necessary truth. (…) » Not so, I reply, for I do not try to imagine *myself* being you, Napoléon, or even simply *myself* being at the battle of Waterloo. This would be impossible indeed, as impossible as imagining you being Napoléon or fighting in that battle. What I can do is imagine being Napoléon, i.e. having the experiences he must have had on the battlefield, or on other occasions. And this, as we have learned, is quite a different thing from imagining myself in some situation or other. (Vendler 1979 : 172-3)

This provides us with another argument against the reflexivist. If we put the self into the content of what is imagined, as the reflexivist does by construing the content of *de se* thoughts as reflexive, we end up with the embarrassing conclusion that what

we imagine when we imagine being Napoléon is something impossible. In the SMR framework the content of basic *de se* thoughts — the sort of thought one entertains when one imagines being Napoléon — is a personal proposition, which does not contain oneself as a constituent but only the properties which one self-ascribes. The subject only comes into the picture in the evaluation phase.

Still we have a problem. Even if the content does not include the self, the self comes into the picture at the next step, in the act of entertaining that content in the relevant mode. If this act is represented as the act of *self-ascribing* the property that is the content of the thought, just as in perception but with an additional element of pretense (in order to distinguish imagination from genuine experience), then don't we have the same contradiction one step later ? How can I (pretend to) self-ascribe the property of being Napoléon and fighting the battle of Waterloo, if those are properties that it is impossible for me to instantiate ?

One response involves distinguishing the possible from the imaginable. Even if the situation is (metaphysically) impossible, still it is imaginable, one might argue. I do not find this response attractive and I will not consider it any further. The correct response, I think, is implicit in the passage just quoted from Vendler : when I imagine being Napoléon and fighting the battle of Waterloo, I imagine certain properties being instantiated, but it is not to myself that I (self-)ascribe the properties in question. Those are properties that I ascribe to *Napoléon*.

Cartesian selves

In 'Imagination and the Self', Bernard Williams discusses what it is for someone to imagine being Napoléon. Three characters seem to be involved : the person doing the imagining (the imaginer), the person who is imagined to be Napoléon (the imaginee), and Napoléon. The crucial question is: who is the imaginee ? Like Vendler, Williams points out that, if the imaginee is the imaginer, we get something metaphysically impossible, for the imaginer could not really be Napoléon. This suggests that the imaginee is *not* the imaginer's actual self but what Williams calls an 'attenuated' self, deprived of all the properties which make it impossible for the actual imaginer to be Napoleon. But what, exactly, is such an attenuated self ?

If we press this hard enough, we readily get the idea that it is not necessary to being *me* [in the attenuated sense] that I should have any of the individuating properties that I do have, this body, these memories, etc. And for some of them, such as the body, we may think that it is not necessary to have one at all… The limiting state of this progress is the Cartesian consciousness : an 'I' without body, past, or character. (Williams 1973 : 41)

At this point, says Williams, we reach an impasse :

Suppose I conceive it possible that I might have been Napoleon — and mean by this that there might have been a world which contained a Napoleon exactly the same as the Napoleon that our world contained, except that he would have been me [in the attenuated sense of 'me']. What could be the difference between the actual Napoleon and the imagined one ? All I have to take to him in the imagined world is a Cartesian centre of consciousness ; and that, the real Napoleon had already. Leibniz, perhaps, made something like this point when he said to one who expressed the wish that he were the King of China, that all he wanted was that he should cease to exist and there should be a King in China. (Williams 1973 : 42)

Williams's way out

To break the deadlock, Williams suggests *doing without the imaginee*. There are, he says, only two characters involved: the imaginer, and Napoléon. It is *to Napoléon himself* that the imagined properties and experiences are ascribed :

Consider… the narration… appropriate to this sort of imagination. It is going to be of the general form : 'I have conquered ; the ideals of the Revolution in my hands are sweeping away the old world. Poor Maria Walewska, I wonder where she is now' and so on and so on, according to whatever knowledge or illusions I possess about Napoleon. Now suppose that we actually heard someone saying things like this. In general, when we hear utterances in the first person, there is only one question to be asked relative to the identity of

164

the 'I' involved : 'Who is the speaker ?' But in the case of utterances as unlikely as this, there are two questions : 'Who is the speaker ?' and 'Who is it that he either believes that he is, or is pretending to be ?' In the present case, the latter alternative is in question : a man engaged in an imaginative narration like this would be a man pretending to be, or playing the rôle of, Napoleon. The 'I' of his discourse is to be taken as an 'I' uttered by Napoleon ; who it stands for, if it stands for anybody, is Napoleon. But of course, this being the playing of a rôle, the actual utterer is someone else, who in the next moment may use 'I' in its ordinary way with respect to his ordinary self. (…) [Similarly,] what I am doing, in fantasy, is something like playing the rôle of Napoleon… In the description of this activity, only two people need figure : the real me and Napoleon. There is no place for a third item, the Cartesian 'I', regarding which I imagine that it might have belonged to Napoleon. (Williams 1973 : 44-45)

On this view, not only I am not a constituent of the state of affairs which is the content of my imagination, but even if we look at the complete content (the 'Austinian' proposition), including the situation in which the imagined state of affairs is imagined to hold, I — the actual imaginer — do not come into the picture. The content of the imagination is assumed to hold in *Napoléon's* situation.

*The quasi-*de se

What I have just said suggests that imagination is not 'just like perception with an additional element of pretense'. There is a much bigger difference between the basic experiential modes (perception and memory) and the imagination.

Memory and perception are reflexive states such that what they represent concerns the rememberer or the perceiver : the rememberer and the perceiver stand in some actual-world relation R to what the memory or the perception represents. Imagination is different : what it represents need not concern the imaginer. While, in perception or memory, I self-ascribe the property of being R-related to what is represented, in imagination no R-relation to the imagined scene is imposed by the act of imagining. The real-world connection between the experiencer and what is experienced is suppressed when we turn to the imagination. As a result, one can

imagine states of affairs to which we can bear no actual-world connection whatsoever.

This feature of imagination makes it similar to quasi-perception and quasi-memory. I will, therefore, coin the term 'quasi-*de se*' to refer to the type of thought one entertains when one imagines, say, being Napoléon. The type of imagining at stake is clearly first personal, yet the imaginer's self is not involved — not even at the 'evaluation' stage. The properties that are imaginatively represented are not ascribed to the subject who imagines them, but to the person whose point of view she espouses.

Chapter 29

Imagination, empathy, and the quasi-*de se*

The Reflexive Constraint

Following Shoemaker, I imagined the alternative modes of 'quasi-perception' and 'quasi-memory', which are like perception and memory, save for the fact that the content of the state need not concern the person in the state. In quasi-perception and quasi-memory, it may not be the person in the experiential state, but some other person, who is *R*-related to what the state represents.

Like Evans, I emphasized how dissimilar quasi-perception and quasi-memory are from ordinary perception and ordinary memory. Their 'semantics', as I put it, is different. Perception and memory obey what we may call the Reflexive Constraint : if the state is veridical, the person undergoing it must be *R*-related to what it represents. The Reflexive Constraint does not apply to quasi-perception and quasi-memory. Now quasi-perception and quasi-memory are counterfactual modes — modes of experience that do not actually exist — but we have just seen that, among the existing modes, one is similar to quasi-perception and quasi-memory in that it does not obey the Reflexive Constraint. That mode is the imagination. The person in whose situation the imagined state of affairs is supposed to hold need not be the imaginer himself ; it may be anybody, including Napoleon, or the last man to be alive on Earth. The imaginer 'sees' the world vicariously, through the eyes of his imaginative target, just as the quasi-rememberer vicariously 'remembers' the experiences of another person.

It follows that we have to make room for a new kind of *de se* thought, or rather, for a distinction within basic *de se* thoughts. Basic *de se* thoughts involve a personal proposition on the side of the *lekton*, i.e. a property of persons, which the subject of the thought self-ascribes in the standard cases in which the Reflexive Constraint holds. But in imagination, as we have just seen, the Reflexive Constraint does not hold. It follows that there are two possibilities, for basic *de se* thoughts entertained in the imaginative mode.

Two types of first person imagination

In one type of case, the imagination concerns the subject himself. For example, the subject imagines (himself) being a racing driver. In such cases, says Williams,

> There is no great problem concerning the *me* that the fantasy is about : it is the actual empirical me, or more or less so. This does not mean, of course, that in order to entertain this fantasy of myself as a champion racing driver I have to engage in an elaborate work of intercalating racing-driving activities hypothetically into my past carreeer, or extending hypothetically my future career so as to embrace them ; I do not have to join the imagined activities in any determinate way on to my actual history. Nevertheless, I am, very often, putting quite a lot of my actual self into it, and where not consciously doing this, am prepared, as it were, to accept a lot of my actual self in the fantasied scene. It is, for example, relative to my real wants, ambitions, and character that the imagined happenings are, to me in them, satisfying or upsetting. (Williams 1973 : 39)

Note that we are talking about first person imagination : the sort of imagination that can be reported by saying 'He imagined being a racing driver', using the PRO construction. What we now see is that, in the case of imagination, such reports can be understood in two ways. What is imagined in the first person way may concern the imaginer's actual self, as we have just seen, or it may concern some other person, as in the Napoleon example. There is no such ambiguity of reports using the PRO construction when the reported state is a state of perception or memory, because the Reflexive Constraint holds for such states. So imagination reports miss a reading which memory and perception reports have (the *de re* reading), but they also enjoy a reading (the 'quasi-*de se*' reading) which perception and memory reports do not have.

As Williams notes (1973 : 44), a quasi-*de se* form of imagination would not be appropriately reported by using the reflexive form 'I imagine myself being Napoléon/fighting the battle of Waterloo'. Only the PRO construction can be used. This is presumably due to the fact that the imaginer's self is not involved at all in quasi-*de se* imagination.

I said that imagination is an experiential mode for which the Reflexive Constraint does not hold. Now it may be denied that imagination is an experiential mode : imagination involves no genuine experience, but only pretense, one may argue. If so, then perhaps the Reflexive Constraint holds of experience generally.

Whatever we think of imagination and its relation to experience, I do not think it can be maintained that the Reflexive Constraint holds of experience generally. There are quasi-*de se* thoughts outside the realm of the imagination. Empathy is a case in point. In empathy, one entertains representations which have the perspectival character that is characteristic of basic *de se* thoughts, but which concern some person distinct from the empathiser.

Suppose I see someone hammering his own finger. In a surge of empathy, I think : 'Ouch, it hurts !' This sentence arguably expresses a personal proposition, true at a person and a time if and only if that person feels pain at that time ; but the relevant person — the person with respect to which the personal proposition is asserted — is not the the person making the assertion (the empathizer), that is, myself, but the person I empathize with. Still, the representation has a clear experiential component : feelings and emotional reactions are involved, which are integral to the mental state, even though the person which the content of the state concerns is not the person in the state. Here, we have an experience that does not obey the Reflexive Constraint.

Book III

EGOCENTRICITY AND BEYOND

Part Nine

Unarticulated constituents in the *lekton* ?

Chapter 30

The context-dependence of the *lekton* :

how far can we go ?

Redefining the lekton

In Chapter 10, following Barwise, I argued that the *lekton* corresponds to the cognitive content of an utterance, while the Austinian proposition corresponds to its complete truth-conditional content, as determined in part by psychological factors and in part by environmental factors. In Chapter 15, however, we saw that the *lekton* can be equated to the utterance's cognitive significance only if it is purely psychological, hence context-independent. Lewis indeed promoted a version of Moderate Relativism where the *lekton* is the meaning of the sentence-type (modeled as a function from context-index pairs to truth-values) : everything nonpsychological — everything not 'in the head' — is rejected on what I call the situational side, that is, the side of that against which the *lekton* is evaluated. This includes (for Lewis) both the context and the index it determines.

At this point the question that arises is this : what reason can we have for insisting that the *lekton* must be context-dependent (in the manner of Kaplan's 'contents'), hence not purely psychological ? At the end of chapter 15 I said that I would provide another justification for construing the *lekton* as context-dependent (i.e., a justification distinct from the cognitive significance justification, which precludes context-dependence). The time has come to actually provide it.

Rather than equate the *lekton*/situation distinction with that between purely psychological and environmental determinants of truth-value, as Lewis does, I suggest that we align it with the distinction between content and mode, presented and elaborated in Book II. This is most natural since, as we have seen, the mode fixes the situation of evaluation, and thereby determines aspects of the complete content that are not articulated in the representation, hence that are not part of its content in the narrow sense. In this framework, as I will now show, the view that the *lekton* is context-dependent can be argued for.

There is a clear distinction between those elements of the complete content that fall on the situation side (determined by the mode), and those elements that are aspects of the content to be evaluated with respect to the relevant situation. Let us consider visual perception once again. Besides the constituents of the visual scene — the various entities that are seen — there are additional elements that are relevant to the veridicality of perception but which do not figure among the elements of the scene. Such elements are, as Perry says, 'unarticulated'. Thus the subject who sees is not herself an aspect of what is seen, but an aspect of the situation with respect to which what is seen is to be evaluated. The subject's perception that *p* is veridical if and only if it is the case that *p in the situation which affects the subject's senses and causes his or her current perceptual experience*. The subject is not a constituent of what is seen, but it features in the complete (Austinian) content of the experience via the situation component, as determined by the mode.[70]

Now when it comes to those elements which *are* represented as part of the visual scene, hence belong to the *lekton*, we have to draw a distinction between their internal representation and their actual identity as determined by external, environmental factors. Going back to Vendler's example, suppose that, from above the cliff, I point to one of the swimmers and say : 'Imagine he cries for help ; what would you do ?' The swimmer I point to is a constituent of the imagined scene. My addressee is invited to imagine *that swimmer* crying for help. Who the swimmer actually is depends upon the context : it depends upon the identity of the person to whom I happen to point. To be sure, that identity does not matter much in the present

---

[70] Bernard Williams makes much of this distinction in his discussion of Berkeley's famous argument, to the effect that we cannot imagine anything unperceived. (To make Berkeley's argument more appealing Williams changes the formulation and discusses the possibility of *visualising* something unperceived.) « Even if we accept the description of visualising as *thinking of oneself seeing* », Williams says, « [that] does not mean that an element or feature of what I visualise is that it is being seen. (…) We can in fact visualise the unseen, because the fact that in visualisation I am as it were seeing is not itself necessarily an element of what is visualised » (Williams 1973 : 34-35).

case. What the addressee is invited to imagine is, in effect, a scene in which *someone with the same properties as those we see the actual swimmer to have* starts crying for help. But if we turn to the basic experiental modes, perception and memory, the situation changes. If we *see* the swimmer crying for help, the content of what we see involves the swimmer — that very individual. He is seen as instantiating certain properties, but his numerical identity is also relevant to what we see : what we see is *that swimmer* (not some other person with the same properties) crying for help.

Of course, it is open to us to say, with Lewis, that in the visual case just as in the imagination case, the *lekton* should only include the purely psychological content of the representation : a representation of a swimmer standing in a certain contextual relation to us, and crying for help. Since the identity of the swimmer is fixed by the situation (as determined by the mode), it should not figure into the content, on the Lewisian view. I grant that this is a perfectly coherent position to take, but there is an alternative position which is no less coherent : we may construe the perceived scene as a state of affairs in the world, with a real individual as a constituent (the swimmer, whose identity depends in part upon the context), and still distinguish that scene and its constituents from the elements that are not constituents of the scene, but only come into the picture via the situation of evaluation as determined by the mode. Those unarticulated constituents (e.g. the perceiver) do not correspond to anything *in* the representation. There is an obvious difference between the perceiver and the swimmer in that respect. The swimmer is visually represented (hence he is a constituent of the *lekton*), but the visual representation by itself does not fix his identity, which depends upon the context (hence the *lekton* is context-dependent). In contrast, the perceiver is not visually represented : he is not a constituent in the visual scene (hence not a constituent of the *lekton*), but comes into the picture via the situation of evaluation which the perceptual mode imposes. To use Perry's terms, the swimmer is 'articulated' in the representation, though in a context-sensitive manner ; the perceiver remains 'unarticulated'.

To conclude, what justifies maintaining the three-level version of the two-stage picture is a pair of important distinctions : that between content and mode, and that between indexicality and unarticulated constituency (a distinction which applies to mental as well as to linguistic representations). To capture those distinctions properly, we need the Austinian proposition with its two components (*lekton*, and

situation), and we need to acknowledge constituents of the *lekton* whose identity depends upon the context, as the reference of indexicals does. The three levels we end up with are : the meaning of the sentence-type (or, in the mental realm, the narrow psychological content of the representation), the context-dependent *lekton* (corresponding to Kaplan's content), and the Austinian proposition including a situation of evaluation in addition to the *lekton*.

<div align="center">

Unarticulated constituents in the lekton *:*

Barwise vs Perry

</div>

Aligning the *lekton*/situation distinction with that between content and mode enables us to make sense of the context-dependence of the *lekton*. But how far can we go in that direction ? How context-dependent can the *lekton* be ? Here, we have a choice between two possible positions. We may hold that only articulated constituents — elements of content to which something in the representation explicitly corresponds — can figure into the *lekton* ; or we may tolerate unarticulated constituents in the *lekton* as well as in the situation. On this issue Barwise and Perry took conflicting stances. For Barwise, the *lekton* is the articulated content of the representation ; for Perry, the *lekton* may well include unarticulated constituents.

Barwise uses the Holmes-Watson example (chapter 10) to illustrate his disagreement with Perry (Barwise 1989a: 240). In the example, Holmes and Watson face each other, with the salt and the pepper standing in between them. Holmes says 'The salt is left of the pepper', because the salt is left of the pepper from Holmes's perspective. From Watson's perspective, the pepper is left of the salt, but Watson is mistaken as to which shaker is which and he wrongly says 'The salt is left of the pepper'. Holmes and Watson apparently 'say the same thing', but Holmes is right and Watson wrong. Some unarticulated constituent must be involved, which accounts for the difference in truth-value. This unarticulated constituent is the *perspective*: the salt is on the left from Holmes's perspective, but it is not on the left from Watson's perspective. (That is why Holmes is right and Watson wrong.) Thus far Barwise and Perry agree, but now a decision has to be made : the unarticulated constituent may be fed into the content to be evaluated (the *lekton*), or into the situation which that content concerns.

On the first option, both Watson and Holmes are talking about the same 'objective' situation (the situation they share), but they state different facts about that situation. The facts they state are, respectively:

Holmes:
Left-of (salt, pepper, perspective H)
Watson:
Left-of (salt, pepper, perspective W)

Watson's and Holmes's perspectives turn out to be (unarticulated) constituents of the facts which they state. According to Barwise, that is the view which Perry favours.

On the second option, taken by Barwise, Holmes and Watson assert the same (relativized) fact:

Left of (salt, pepper)

However, Holmes and Watson talk about *different situations*. The situations are individuated in terms of Holmes's and Watson's subjective perspectives on them.

Whatever we think of such 'perspectival situations', and the specific problems they raise,[71] there is a more fundamental issue at stake. According to Barwise, writing in the late eighties, Perry *generally* treats 'unarticulated constituents' as constituents of the content to be evaluated, rather than aspects of the situation with respect to which the content is evaluated (the situation which the representation 'concerns', in the terminology of Perry's early paper, 'Thought without representation'). If this is right, then Perry must have changed his mind since that early paper. For the paper's main point was that more cases of unarticulatedness can be handled in terms of the 'concerning' relation than one might as first suppose. Be that as it may, it is important to realize that, even in 'Thought without representation',

---

[71] The main problem perspectival situations raise is that they are not 'objective' enough to play the role of situation; or so it may be thought. That difficulty can (perhaps) be overcome by viewing perspectival situations as 'reflexive situations' which include a viewer with a perspective *on* the situation. Be that as it may, I am not specifically concerned with perspectival situations in this chapter. I use them only for illustrative purposes.

Perry remained very cautious and resisted the sort of generalization which characterizes Barwise's approach. Not all instances of unarticulatedness, he then suggested, can be handled in terms of the concerning relation. In the next chapter, I will present what I take to have been Perry's criterion, at the time of 'Thought without representation', for picking out the cases that can be so handled, and for maintaining that in other cases, an unarticulated constituent *is* a constituent of the *lekton* rather than an aspect of the situation which the representation concerns.

Chapter 31

Unarticulatedness and the 'concerning' relation

Thought without representation

The SMR framework owes much to Perry's pioneering paper, 'Thought without representation', where he introduces the notion of an unarticulated constituent and the distinction between 'concerning' and 'being about'. That distinction comes out most clearly in the case of the Z-landers, a small group of people who « do not travel to, or communicate with residents of, other places » (Perry 1986b/1993b : 212) and have no name for Z-land, the place where they live. As Perry points out, Z-land is an *unarticulated* constituent of the content expressed by the Z-lander's utterance 'It's raining'. The utterance is true if and only if it is raining *in Z-land* ; but the Z-landers do not have a concept or idea of Z-land as opposed to other places. Their weather thoughts 'concern' Z-land, not by virtue of containing a representation of Z-land (in which case they would be 'about' Z-land), but by virtue of their *being in* Z-land. The unarticulated constituent is unarticulated not only linguistically but also mentally: it's a constituent of content directly provided by the environment.

In such cases the mental representation, considered in abstraction from the environment which it concerns, expresses less than a complete proposition. The Z-landers think 'It is raining': the content thus articulated is not fully propositional — it is a propositional function, which is truth-evaluable only with respect to a particular place (determined by the environment). Now, as Perry pleasantly says, « there is a little of the Z-lander in the most well-traveled of us » (Perry 1986b/1993b: 216). The difference between the Z-landers and us is that we do have a notion of the place where we live, as opposed to other places; so we are capable of entertaining a thought *about* the place where we are, such as 'It's raining in Paris, but not in Saint Tropez'. Perry's point, however, is that when we are in Paris (or Palo Alto) and we say or think 'It is raining', we need not think *reflectively* about the place we are in. We can think 'It is raining' and *let the place we are in complete the content of our thought.*

> When I look outside and see rain and grab an umbrella or go back to bed, a relatively true belief, concerning my present surroundings, will do as well as a more articulated one, about my present surrounding. (Perry 1986b/1993b : 216 ; emphasis mine)

Unarticulatedness without relativization

Not *all* instances of unarticulatedness can be handled in terms of 'propositional functions' and the concerning relation, however. Contrary to what many commentators assume, the two distinctions which Perry introduces in 'Thought without Representation' (concerning/being about, and articulated/unarticulated) are not coextensive. For Perry, there *are* unarticulated constituents that are genuine constituents of content, rather than aspects of the situation with respect to which the content is evaluated.[72] The following example, discussed by Perry, presumably falls into that category :

> Suppose… that my son has just talked to my older son in Murdock on the telephone, and is responding to my question, « How are things there ? » Then his remark ['It is raining'] would not be about Palo Alto [the place where he is], but about Murdock... My son belief [is] about Murdock, and his intention [is] to induce a belief in me that [is] about Murdock by saying something about Murdock. Here it is natural to think that we are explaining which unarticulated constituent a statement is about, in terms of something like the *articulated* constituents of the beliefs and intentions it expresses. (Perry 1986b : 211).

---

[72] Sometimes, Perry even speaks as if *all* unarticulated constituents were constituents of content rather than aspects of the situation which the representation concerns. « [A] statement is *about* the unarticulated constituent[s], as well as the articulated ones » (Perry, Thought without Representation, 1986b/1993b : 209). Barwise 1989a understands Perry's unarticulated constituents that way. On this usage — which I will not follow — Z-land is only a *prima facie* unarticulated constituent of Z-landers' thoughts (but in fact, not an articulated constituent at all, since it is not a constituent of content in the first place).

In this case at least, according to Perry, the statement is 'about' the unarticulated location ; it does not (merely) 'concern' it. The content of the assertion is a classical proposition with the location as (unarticulated) constituent, rather than a propositional function evaluated with respect to that location.

The Externality Principle

Though he did not discuss the issue explicitly, the following passage seems to me representative of Perry's view, in 'Thought without representation', regarding the cases which can and those which cannot be handled by appealing to relativized propositions and the concerning relation:

> In those parts of our life where there is an *external guarantee* that the weather information we receive and our actions will concern our own locale, there is no reason for our beliefs to play the internal coordinating role they need to at other times. (Perry 1986b : 216 ; emphasis mine)

This actually covers two sorts of case. There are, on the one hand, the cases in which the subject has no representation whatsoever of the relevant parameter, which only the theorist can articulate. That is the 'Z-lander' sort of case. There are also the cases in which the subject herself can articulate the relevant parameter, but need not do so because the value of the parameter is fixed by the environment, without any need for the subject herself to *cognitively discriminate* the situation of concern from other possible situations. That is what happens in the mode of thinking or discourse that specifically concerns local weather :

> Those belief states that directly control behavior for local weather merely *concern* local weather, rather than being about it. All believers who had just seen rain and were about to open their umbrellas [should] be reckoned as believing the same propositional function, but the truth conditions of their beliefs... differ with their location. (Perry 1986b : 217)

So 'It is raining' expresses a propositional function when it is uttered in talking about local weather. Even though 'It is raining', in such circumstances, turns out to have the same truth-conditions as 'It is raining here', they are not synonymous : 'It is raining' expresses a place-relative propositional function, while 'It is raining here' articulates the place which therefore goes into the evaluated content instead of being simply part of the circumstance of evaluation. In the 'It is raining' case, it is the mode of discourse, as Perry says, that constrains the representation to be evaluated with respect to the place in which it is tokened. The content itself is locationally neutral — it does not specify the place of rain.

The cases that presumably cannot be handled in this way, according to Perry, are the cases in which it is *incumbent upon the subject to discriminate what her thought or statement is tacitly about*, because there are several possible options and no 'external fact' to pick out one. That is what happens in the mode of thinking or discourse about non-local weather, as in the Murdock example. In the Murdock case, according to Perry, the subject expresses the belief that it's raining *in Murdock*. The place (Murdock) is not linguistically articulated, but it is a constituent of content nonetheless.

All this suggests that Perry's criterion for treating an unarticulated constituent as an aspect of the situation rather than as an implicit ingredient of the *lekton* may have been the following principle :

> *Externality Principle*
> For an unarticulated constituent to be an aspect of the situation rather than an implicit ingredient of the *lekton*, it must be contributed by the external environment rather than cognitively discriminated.

Whether or not Perry actually endorsed something like this principle, I reject it and will criticize it in the next two chapters. I will ultimately argue against the very idea that there are unarticulated constituents in the *lekton* : I will defend Barwise's position, according to which the *lekton* is the *articulated* content of the representation (and everything unarticulated falls on the situational side).

Chapter 32

Three (alleged) arguments for the Externality Principle

The mental articulation argument

The Externality Principle may be tentatively justified by means of three arguments. The first one, suggested by Perry's description of the Murdock case, runs as follows. Perry insists that, in that case, it is not the location of the speaker, but his *intentions* and *beliefs*, which determine the place on which the truth-value of the statement depends. Since that is so, the place in question has got to be mentally represented, that is, articulated in the mental representations which underlie the speaker's utterance. In particular it must be articulated in the belief which the utterance expresses (the belief that it is raining in Murdock). So Murdock is an *articulated* constituent of the content of the belief.[73] If, as seems legitimate, we assume that the content of the utterance is the same as that of the belief which it expresses, it follows that the place is a constituent of the content of the utterance as well, even if it is not articulated in the utterance itself (but only in the belief that the utterance expresses). Conclusion : the place is an articulated constituent of the content of the subject's belief — the subject believes, and intends to communicate, that it is raining *in Murdock* — and an unarticulated constituent of the content of the utterance ('It is raining') which expresses that belief. It is *not* simply an aspect of the situation which the utterance concerns.

I do not accept this piece of reasoning, and I reject its conclusion. I have no quarrel with the assumption that the content of an utterance is the same as that of the belief which it expresses,[74] nor with the premiss that Murdock must be mentally represented if the speaker is to be credited with the appropriate communicative intentions. Still, there is a premiss in the above reasoning which I find unpalatable.

---

[73] « Here it is natural to think that we are explaining which unarticulated constituent a statement is about, in terms of something like the *articulated* constituents of the beliefs and intentions it expresses. » (Perry 1986b : 211)

[74] This assumption corresponds to my 'Congruence Principle' (Recanati 1993 : 54-55).

The fact that Perry's son must think of Murdock and intend to say something about Murdock when he utters 'It is raining' possibly entails that Murdock is articulated in some mental representation of his, but does *not* entail that the belief he expresses by his utterance 'It is raining' is the locus of that articulation. So what I deny is the legitimacy of the inferential step corresponding to the phrase 'in particular' in the following portion of the argument :

> The place in question has got to be mentally represented, that is, articulated in the mental representations which underlie the speaker's utterance. *In particular* it must be articulated in the belief which the utterance expresses (the belief that it is raining in Murdock).

I grant that the relevant place (Murdock) is articulated in some mental representation in the speaker's mind, which mental representation crucially serves as background for his utterance 'It is raining', but I deny that the mental representation in question must be identical to the very belief he expresses by this utterance. Without begging the question, we cannot rule out the following option : (1) by his utterance the speaker expresses the belief that it's raining ; (2) that belief is location-relative and *concerns* Murdock ; (3) what determines that the belief concerns Murdock (rather than some other place) is some other mental representation in the speaker's mind, involving Murdock as an articulated constituent.

In general, the contextual facts which fix the value of the situational parameter for a given mental representation may well be cognitive factors, involving *other* mental representations. To take an example I have used elsewhere (Recanati 1997, 2000b), suppose I say : « Berkeley is a nice place. There are bookstores and coffee shops at every corner. » This is a two-sentence discourse. Berkeley is an articulated constituent of the first statement, and an unarticulated constituent of the second statement. Nothing prevents us from saying that the second statement *concerns* the place which the first sentence explicitly mentions. The fact that that place is cognitively discriminated via the mental representation corresponding to the *first* sentence does not entail that it is articulated also in the mental representation corresponding to the *second* sentence. On the contrary, the fact that the subject has just entertained a representation explicitly about Berkeley contributes to explaining why the second representation concerns that city. Likewise, I think it is the mental

representation corresponding to Perry's question 'How are things there ?', not that corresponding to his son's answer 'It is raining', which articulates Murdock.[75] The place thus articulated in the question can serve as the situation which the answer concerns.

Given all this, I see no reason not to accept that, in many cases (including the Murdock case), the situation which an utterance or thought concerns is determined by cognitive factors such as the topic of the conversation or what the thinker is mentally focussing on, rather than by external facts like the location of the speaker. In such cases, admittedly, the situation s which the representation r concerns will itself have to be somehow represented or articulated — it will have to be cognitively discriminated — but that would raise a problem only if that entailed that s is articulated *in r*. As we have just seen, that consequence does not follow.

The behavioural argument

The second argument in favour of the Externality principle is the following.[76] One reason to hold that two persons who look out the window and say 'It's raining' believe the same thing (even though their locations are different) is that, because of their belief, they are disposed to act in the same way : by grabbing their umbrellas, cancelling plans for picnic, etc. It follows that what they believe — in one sense of that phrase — can be construed as a place-relative proposition, or propositional function. But when the unarticulated location of the raining event is not fixed by the environment (i.e. by the place the believer is in), but by cognitive factors — as in the Murdock case — the relevant dispositions are absent. After his phone call to his

---

[75] It might be argued that the answer is actually elliptical for 'it's raining there' (just as, in answer to 'what is Jones doing ?', an utterance of 'Raking the leaves' would be elliptical for 'Jones is raking the leaves'). If this is right, then Murdock is not an unarticulated constituent at all. I assume, with Perry, that there is no ellipsis in the linguistic sense and that Murdock is an unarticulated constituent : the issue between us is whether the unarticulated constituent in question goes into the content or merely features as an aspect of the situation the utterance concerns.

[76] I am indebted to Michael Devitt for suggesting that argument in discussion during the 4th Barcelona Workshop on the theory of reference.

brother, Perry's son knows that it is raining in Murdock, but, being himself in Palo Alto, he is not disposed to grab his umbrella. It seems, therefore, that Perry's son does *not* believe the same thing, or does not express the same belief, when he says 'It is raining' after the phone call with Perry's older son, as he expresses when he says 'It is raining' after looking out the window in Palo Alto. In the latter case it makes sense to say that he believes the propositional function that it is raining, hence harbors the same belief as anyone who looks out the window and sees rain ; but in the Murdock case what Perry's son believes is the complete proposition that it is raining *in Murdock*. That explains why Perry's son, in the Murdock case, is not disposed to act in the way one acts when one believes the propositional function that it is raining : by grabbing one's umbrella and cancelling out picnic plans.

I will discuss that argument in Chapter 41, for it connects up with issues of simulation and shiftability to be dealt with then. I now turn to the third potential argument in support of the Externality Principle : the argument from invariance, which I will discuss and reject in Chapter 33.

The argument from invariance

The argument I have in mind is implicit in the following passage from Perry :

> In cases in which *the same unrepresented parameter* is relevant to a whole mode of thinking or discourse, we should classify each specific belief or utterance with a propositional function. The truth-value would be that of the proposition obtained by applying the function to the value of the parameter fixed by facts about the whole system. (Perry 1993a: 221 ; emphasis mine)

Here what matters is the *invariance* of the unarticulated constituent. This condition of invariance is trivially satisfied in the Z-land case :

> All of the information (or misinformation) [the Z-landers] get about the weather, through observations or reports of others, is about Z-land. All of the actions they perform, in light of their weather beliefs, take place in Z-land, and are appropriate or not depending on the weather there. (Perry 1986b/1993b : 213)

The invariance condition is also, though less obviously, satisfied in the case of Perry's son looking out the window and saying 'It's raining'. To be sure, in contrast to the Z-land case, the place of rain is liable to vary : each utterance of 'It is raining' concerns the place where the utterance is made, and that place can vary. But in another sense, the place does not vary : it is always *the place where the subject is*. The parameter is fixed, even if its values vary as a function of speaker's location. In contrast, when the unarticulated constituent is not fixed by the environment, but by cognitive factors, it can freely vary according to the speaker's intentions : the invariance condition is not satisfied. Or at least, that is what the argument from invariance says.

The following passage, which summarizes Perry's view, stresses invariance as the crucial consideration :

> Sometimes all of the facts we deal with involving a certain *n*-ary relation involve *the same object* occupying one of the argument roles. In that case, we don't need to worry about that argument role ; we don't need to keep track of its occupant, *because it never changes*. We can, so to speak, pack it into the relation. For centuries people in Europe assumed that 'being a summer month' was a property of months. July was a summer month, December was not. Once they started to visit the Southern Hemisphere, they had to take account of the relativity to places. July was a summer month in the Norhtern Hemisphere, but not in the Southern Hemisphere. A child who is unconcerned about and even unaware of the weather anywhere but where he is, can treat the issue of whether it is raining or not as a property of a time, rather than a relation between times and places. He says, « It is raining now » rather than « It is raining here now. » (In this case the argument role is not always occupied by the same place, but always occupied by a place with a fixed relation to the agent, the place he is at.) (Perry 2000 : 328 ; emphasis mine)

The same idea can be found in Barwise's paper 'Situations, facts, and true propositions'. Given a situation *s* and a relation *R* with *n* argument roles, Barwise says, an argument role *r* of *R* is *undiscriminated* in *s* iff there is an entity c such that c fulfils that role in any state of affairs of *s* involving *R*. (For an argument role to be undiscriminated is, therefore, for its filler to be invariant.) When that is so, one can

'project' the relation R to a relation R_r of one fewer roles, by suppressing the role in question. This operation Barwise calls 'relativization'. It is justified by the fact that « from the perspective of an agent with focus situation s, the role r might as well not exist, since it is completely undetectable in that situation » (Barwise 1989a : 253). [77]

To sum up, invariance is what justifies taking an unarticulated constituent out of the *lekton*. Externality matters only insofar as it is what guarantees invariance.

---

[77] The dissymmetry, emphasized by Evans, between worlds and times, and the correlative acceptability of relativization-to-worlds as opposed to relativization-to-times (Chapter 2), may itself be construed in terms of invariance. For Evans, there is one and only one world that is actual ; and all our assertions are about that world. In contrast, there are many different times, none of which has a unique status corresponding to that of the actual world. The invariance condition is satisfied in one case, not in the other ; and relativization is legitimate only when that condition is satisfied.

Chapter 33

Invariance

Millikan vs Perry

Ruth Millikan is another author who insists on invariance as the crucial condition determining whether or not an aspect of truth-condition should be treated as an aspect of content narrowly construed. According to her,

> The meaning of the sign is determined as a function of values of significant variables or determinables exhibited by the sign. Put another way, the meaning varies systematically to parallel significant (mathematical) transformations of the sign. (Millikan 2004 : 48)

It follows that

> Aspects of a truth condition are explicitly represented only when expressed as values of variables that can accept alternative values. What is constant over all possible representations in a system often represents quite a lot, but represents it only implicitly. (Millikan 2006 ; emphasis mine)

This criterion enables us to distinguish, among aspects of a truth condition, what is explicitly represented and what is not. She gives the following example :

> Consider… the dance of the honey bee. By its orientation on the hive wall it shows the direction of nectar relative to a line from the hive to the sun. Change the dance orientation and the represented direction of nectar changes accordingly. Direction is explicitly represented in the dance (*by* direction). The bee dance also talks about nectar, but no change will make it talk about something else, say, the direction of milk or peanut butter. Nectar is represented only implicitly. Likewise the hive and the sun. Though all three,

nectar, hive and sun, figure in the dance's truth conditions, none is represented explicitly. (Millikan 2006 ; emphasis mine)

Although Perry and Millikan both appeal to invariance, their views of the matter differ strikingly. As we have seen, Perry accepts cases of parametric variation as cases of invariance. The child who says 'It's raining' talks about the place where he is. « In this case », Perry says, « the argument role is not always occupied by the same place, but always occupied by a place with a fixed relation to the agent, the place he is at. » *Given* the context in which the utterance takes place, the occupant of the argument role is fixed (by the environment) and cannot vary. But for Millikan, this is not enough. She contrasts the dance of the honey bee with the STOP sign (Millikan, p.c.). The location of the STOP sign indicates the location of the stopping event, but the location of the STOP sign is liable to vary, hence the location of the stopping event, which varies accordingly, is an aspect of the content of the sign. It is, in her terms, explicitly represented. To determine what is explicitly represented, what we must do, according to Millikan, is *look and see what aspects of the representation can be transformed in a systematic way to yield other representations in the same system that represent other things*. The location of the STOP sign is one such aspect that gives rise to meaningful transformations, hence is a meaningful component of the overall sign, just like the direction in the bee dance case. Now the same analysis applies to 'It is raining'. The place where the statement is made or the belief held determines the place where rain is said or believed to occur, thus the location goes into the content of the statement/belief and counts as 'explicitly represented', for Millikan.

To be sure, the place of rain is *linguistically* unarticulated in 'It is raining' : no word or morpheme stands for it in the sentence. But for Millikan, that does not mean that the place of rain is unarticulated *tout court*. There is something in the overall sign that stands for the place of rain, she argues, but it is not a linguistic symbol : it is *the place of utterance*. The place of utterance is an element of the sign that stands for itself, that is, contributes itself (that very location) to the content of the utterance : an utterance of 'It is raining' occurring at a certain location says that it is raining *at that location*. The place which is a constituent of the proposition expressed by the utterance is the semantic value of… the place of utterance construed as a

meaningful component of the utterance.[78] Rather than 'unarticulated', the location of the raining event is *reflexively* articulated, Millikan says — it is articulated by means of what she calls a 'reflexive sign' :

> Times that represent times and places that represent places are perfectly ordinary ingredients of natural signs… We can invent a special term for the case where a sign element represents itself. Call these sign elements « reflexive ». (Millikan 2004 : 49)

In the case of the Z-landers, Millikan would agree with Perry that the place is not represented, since it is truly invariant. But she would reject Perry's contention that 'there is a little of the Z-lander in the most well-traveled of us'. Likewise, she agrees with Perry that the subject of perception is not represented, for it cannot vary. The content of perception is selfless — the self comes into the picture only through the architecture of the system. But she treats the place and time of perception differently : *they* are elements of the representation that stand for themselves and contribute themselves to the content of the perception.[79]

Implicit representation

Millikan's notion of invariance is absolute, Perry's notion relative. Being absolute, Millikan's notion of invariance is stronger than Perry's, but Perry's weaker notion is no less legitimate and can serve to ground a notion of implicit content distinct from, and more inclusive than, the notion Millikan defines in terms of absolute invariance.

---

[78] As Frege pointed out, the complete sign involves more than the sentence, it also includes the time of utterance (and, according to Millikan, the place as well).

[79] Note that the time of a remembered event is not *reflexively* represented by the time of the memory episode : we remember an event *after* its occurrence. This suggests that Millikan's theory of reflexive signs is an overgeneralization based on a special case. Rather than granting this, however, she broadens the notion of a reflexive sign so as to encompass what she calls 'relative reflexives' (Millikan 2004 : 53). For a discussion of Millikan's theory of signs, see Recanati forthcoming c.

Let us grant that changes in the time or place of perception are 'transformations' that yield systematic changes in the perception's truth-conditions. Still, there is a sense in which a perception *can only represent what is going on at the time and place of perception*. That is fixed by the perceptual mode. In the same way, an episodic memory *can only represent a past event which the subject has experienced*. That is fixed by the memory mode. As Perry rightly points out, there is an element of invariance here. To be sure, there is also an element of variability. What is fixed by the perception mode is the *relation* which, say, the time of the represented event must bear to the time of the perceptual episode ; and that time can vary. But the relevant relation to the perceptual or memory episode — the relevant 'parameter', as Perry says — is fixed by the mode and cannot vary. This is sufficient to distinguish the values of the parameter in question, which are provided automatically (as a function of the time and place of the perceptual or memory episode), from the elements of content that are explicitly represented and can vary freely, without being constrained or determined in this way.

I conclude that Perry's view and Millikan's do not really conflict — they talk about different things, despite using the same terminology ('implicit' vs 'explicit representation', etc.). Perry and Millikan both accept that a representation has truth-conditions, and each of them draws a distinction — a different distinction — between aspects of the truth-conditions in question. I will assume that the representation's truth-conditions are a certain state of affairs that it represents. Thus the Z-lander's thought 'It is raining' represents the occurrence of rain in Z-land at the time of the thought episode. Millikan's criterion can be used to distinguish, among the various aspects of what is globally represented, what belongs to the *meaning* of the representation and what does not. By her criterion, the meaning of a sign is « determined as a function of values of significant variables or determinables exhibited by the sign ». Both the weather condition (rain) and the time of the raining event satisfy that condition and therefore belong to the meaning of the representation, but the place of rain (Z-land) does not because it is invariant in the absolute sense. Using the weaker notion of relative or parametric invariance as a criterion, Perry draws a different distinction between what is explicitly represented and what is only implicitly represented. What is implicitly represented covers all the aspects of what is represented that are fixed (either absolutely *or relatively*) and cannot vary freely. Presumably, what is explicitly represented (in language or

thought) can only belong to the representation's meaning, in Millikan's sense, since it is compositionally determined ; but not every aspect of the representation's meaning, in Millikan's sense, is explicitly represented in Perry's sense.

To sum up, starting from the representation's overall truth-conditions, Millikan's criterion distinguishes between what belongs to the meaning of the representation and what does not, while Perry's criterion distinguishes between what is explicitly represented and what is not. The distinctions are far from identical : what is explicitly represented (in Perry's sense) does not exhaust the meaning of the representation in Millikan's sense ; nor do the aspects of truth-conditions that are external to meaning exhaust what is implicitly represented in Perry's sense.

What is explicitly represented, in Perry's sense, is what I call the *lekton*. In my framework as in Perry's, the place of rain is only implicitly represented when an ordinary subject looks out the window and says 'It is raining'. The explicit content of the representation is a propositional function — a function from places to truth-values ; and the place which serves as argument to the function is determined via a certain relation (identity) to the place of utterance — a relation which itself is fixed by what Perry calls 'the mode of thinking or discourse'.

The invariance argument rebutted

We have just seen that, in a standard use of 'It is raining' to talk about the local weather, the external environment (viz. the place and time of the utterance event) determines the place or time which the utterance implicitly represents. But what about the case in which the place of the raining event is determined by cognitive rather than environmental factors ? What about the Murdock case ? Is the place explicitly or implicitly represented, in such a case ?

If we accept the Externality Principle, as it seems that Perry does, we should say that the location of the raining event is explicitly represented and belongs to the *lekton*, despite the fact that it is linguistically unarticulated. But why should we accept the Externality Principle ? As we have noticed, what really matters is invariance. According to the argument from invariance, we should accept the Externality Principle because « when the unarticulated constituent is not fixed by the environment, but by cognitive factors, it can freely vary according to the speaker's intentions : the invariance condition is not satisfied. » If the place of rain is not the

place of utterance, but is fixed by what the speaker has in mind or intends to refer to, it can vary indefinitely. In one context the speaker will be talking about Murdock and his utterance 'it is raining' will say that it is raining in Murdock, in another context he will be talking about Paris. There is limitless variability here. Hence the place is explicitly represented despite being linguistically unarticulated.

This argument in favour of the Externality Principle I find very weak. Once we have admitted the legitimacy of parametric invariance, one can easily extend that notion so as to cover the Murdock case as well. Perry himself distinguishes several 'modes of thinking or discourse'. He distinguishes *talk about the local weather* from *talk about the weather at other places*. When one talks about the local weather, the relevant parameter is the place of utterance : one's weather talk concerns the place where one is. What is explicitly represented is the weather — the place is automatically determined by external facts regarding the utterance. Why not similarly construe the other type of discourse as a language game or discourse mode that sets up another parameter ? While a rain statement in the 'local-weather-talk' mode concerns the place of the statement, a rain statement in the other mode concerns whichever place happens to be cognitively salient or is currently the topic of conversation. Clearly, what determines the value of *that* parameter is not the external environment, but cognitive factors. Still, the notion of parametric invariance applies : the mode of discourse fixes the relevant parameter (the place of utterance *vs* the place one is talking about), and the value of the parameter is determined by the relevant sort of factors ('external' or 'cognitive', as the case may be). There is variation in both cases, and in both cases the variation is constrained by the relevant parameter.[80]

I conclude that neither the mental articulation argument, nor the argument from invariance, forces us to accept the Externality Principle. Since the Externality Principle is the only reason I can think of for insisting that some unarticulated constituents (those that are cognitively rather than externally determined) must belong to the *lekton*, and there is no compelling reason to accept the Principle, I will continue to assume that the *lekton* only contains elements that are explicitly articulated.

---

[80] This parametric approach to nonlocal weather talk will be elaborated in Chapters 40-41.

Part Ten

Self-relative thoughts

Chapter 34

The problem of the essential indexical

In Chapters 14-15 I mentioned two competing explanations of the phenomenon of indexical belief. Lewis accounts for indexical belief by claiming that its content is a relativized proposition (a property which the believer self-ascribes). Perry, instead, appeals to the distinction between the object of belief (a classical proposition) and the mode of presentation under which it is believed, and to the correlative distinction between belief and acceptance.

I am going to argue that we need both relativized propositions and modes of presentation in a complete account of the phenomenon — an account that does not ignore the crucial distinction between basic and nonbasic *de se* thoughts. I start with a general statement of the problem of indexical belief, also known (since Perry's celebrated paper) as the 'problem of the essential indexical'. In the next chapter I will present Perry's objection to solutions based on relativized propositions. As we will see, there is a close connection between Perry's argument against relativized propositions and the Externality Principle which we have just discussed (and rejected).

The eternalization thesis

In *The Logical Syntax of Language*, Carnap said he was dealing "only with languages which contain no expressions dependent upon extra-linguistic factors" (Carnap 1937 : 168). Carnap's disciple Bar-Hillel lamented that this "restricts highly the immediate applicability" of Carnap's views to natural languages since "the overwhelming majority of the sentences in these languages are *indexical*, i.e. dependent upon extra-linguistic factors" (Bar-Hillel 1963 : 123). Bar-Hillel ventured the hypothesis that "more than 90 per cent of the declarative sentence-tokens we produce during our life-time are indexical sentences and not statements" (Bar-Hillel

1954 : 76 ; a 'statement', in his terminology, is a sentence that expresses the same proposition whichever context it occurs in).

Despite his emphasis on the pervasiveness of indexicality, Bar-Hillel accepted that "a judgment [i.e. an ordered pair consisting of a sentence and a context] with an indexical sentence as first component can always, without loss of information, be transformed into a judgment with a statement as a first component, keeping the second component intact" (Bar-Hillel 1954 : 76). Thus if, in context *c*, John says 'I am hungry' and thereby expresses the proposition that John is hungry at *t* (the time of *c*), he can express the same proposition in the same context by uttering "John is hungry at *t*".[81]

The thesis that indexical sentences can always be rephrased into a context-invariant form without loss of information deserves a name. Let us call it the 'eternalization thesis'.[82] It used to be very commonly accepted until fairly recently. In the late sixties a general principle — the principle of 'Expressibility' (Searle) or 'Effability' (Katz) — was put forward, which entails the eternalization thesis as a special case. According to that general principle, *whatever may be conveyed by uttering a sentence S in a context c can also be literally expressed, in a context-independent manner, by means of a fully explicit sentence S'*. One consequence of the principle is that "cases where the speaker does not say exactly what he means — the principal kinds of cases of which are nonliteralness, vagueness, ambiguity, and incompleteness — are not theoretically essential to linguistic communication" (Searle 1969: 20). Indexicality also counts as theoretically dispensable.[83] In principle, we can

---

[81] Bar-Hillel follows Carnap here:

> The logical character of [nonindexical sentences] is... invariant in relation to spatio-temporal displacements; two sentences of the same wording will have the same character independently of where, when, and by whom they are spoken. In the case of [indexical sentences], this invariance can be attained by means of the addition of person-, place-, and time-designations. (Carnap 1937 : 168)

[82] See Recanati 2004 and 2005a on 'Eternalism' as a general philosophical position regarding the place of context-sensitivity in natural language.

[83] At least this follows from the Principle of Effability as formulated by Katz. Searle's formulations are not as clear-cut. On the relations between the two principles, see my article 'The Limits of Expressibility' (Recanati 2003).

always replace an indexical expression by a nonindexical one. Instead of saying 'Thank God, he's gone' I can say 'The man who just asked the stupid question about the relation betwen the mental and the physical has, thank God, left the room' (Katz 1977: 20) ; and instead of saying 'That man is a foreigner' I can say 'There is one and only one man on the speaker's left by the window in the field of vision of the speaker and the hearer, and he is a foreigner' (Searle 1969: 92). To be sure, that way of speaking would not be very convenient in practice. As Katz puts it, indexicality "allows speakers to make use of contextual features to speak far more concisely than otherwise" (Katz 1977: 19).

A pragmatic constraint on indexical-elimination

In practice, the transformation of indexical sentences into a context-invariant form raises « formidable problems » (Bar-Hillel 1963 :123). The eternalization thesis says that, for any sentence S, context *c*, and proposition *p* which S expresses in *c*, there is a sentence S' such that in every context (including *c*) S' will express that same proposition *p*. In other words, *c* and *p* remaining constant, it is always possible to replace S by a nonindexical sentence S': that is the gist of the eternalization thesis. Still, Bar-Hillel pointed out, there is a sense in which S *cannot* be replaced by S'. Consider a very simple example: the replacement of 'I am hungry' (S) by 'John is hungry at *t*' (S'). S and S' express the same proposition (that John is hungry at *t*) in every context in which John is the speaker and *t* is the time of utterance; but that does not mean that S and S' can be freely interchanged in all such contexts. If the users do not *know* that John is the speaker and *t* the time of utterance, the sentences S and S' will not be taken to express the same proposition, hence they will not be intersubstitutable *in the communicative situation*. S will be actually replaceable by S' only in a small subset of the above set of contexts, namely the contexts in which (i) John is the speaker and *t* the time of utterance, and (ii) the language users are aware of that fact.

In general there is a pragmatic constraint on the transformation from indexical to nonindexical: the language users must *know* the relevant facts in virtue of which S and S' express the same proposition. But it is far from obvious that this constraint *can* be satisfied if the transformation from indexical to nonindexical is to be complete. In the examples I gave above ('The man who just asked the stupid question about the

relation betwen the mental and the physical has, thank God, left the room', 'There is one and only one man on the speaker's left by the window in the field of vision of the speaker and the hearer, and he is a foreigner') the transformation was clearly not complete: there remained various sources of indexicality in the substitution sentences. Arguably, if we try to get rid of *all* indexicals, we will be in a position to do so only by invoking facts which are *not* known to the language users, that is, by violating the pragmatic constraint.

That difficulty, and the pragmatic constraint on which it is based, can be dismissed as irrelevant. Thus Goodman wrote:

> Against such translations, it is sometimes urged that they do not really convey the content of the originals. A spoken "Randy is running now" tells us that the action takes place at the very moment of speaking, while a "Randy runs [tenseless] on October 17, 1948, at 10 p.m., E.S.T." does not tell us that the action takes place simultaneously with either utterance unless we know in addition that the time of the utterance is October 17, 1948 at 10 p.m. E.S.T. Since — the argument runs — we recognize the tenseless sentence as a translation of the tensed one only in the light of outside knowledge, we have here no genuine translation at all. But this seems to me no more cogent than would be the parallel argument that "L'Angleterre" is not a genuine translation of "England" because we recognize it as a translation only if we know that l'Angleterre is England. (Goodman 1951: 268-9)

Goodman's quotation makes clear what is at issue: there are aspects of the intuitive 'content' of the original that are left aside in the nonindexical translation, and at the same time 'outside knowledge' — i.e. information which is not part of that intuitive 'content' — is exploited in producing the nonindexical translation. How is that intuitive notion of the 'content' of the original utterance, what it 'tells us', related to that of the 'proposition' which it expresses, and which the nonindexical translation is taken also to express? Can we discard the intuitive difference in content between the original and its nonindexical translation, as Goodman suggests, on the grounds that they express the same proposition, much as 'l'Angleterre' and 'England' denote the same country?

When we say that 'I am hungry' and 'John is hungry at *t*' 'express the same proposition' with respect to a context *c* in which John is the speaker and *t* the time of utterance, we mean that they *have the same truth-conditions.* Both are true iff John is hungry at *t*. This is captured by saying that they express the same 'singular proposition', consisting of John, the time *t*, and the two-place relation of being hungry at a time. But if we have in mind more fine-grained propositions of the sort Frege was concerned with (what he called 'thoughts'), then it is unlikely that those utterances express the same proposition, even if they have the same truth-conditions. As far as Fregean thoughts are concerned considerations of 'cognitive significance' play a crucial role alongside truth-conditional considerations.

Let us assume that John is rational. At *t* he may well assert 'I am hungry' while, at the same time, dissenting from 'John is hungry at *t*'. (That is possible if, lacking the relevant 'outside knowledge', he does not know who he is, or what time it is. For example, he may mistakenly believe that he is Peter, and that the current time is *t'*.) By Fregean standards, the fact that that is possible shows that the two sentences do *not* express the same 'thought', even with respect to a context in which John is the speaker and *t* the time of utterance. Following this line of argument, it can be shown that *no transformation from indexical to nonindexical is possible without affecting the cognitive significance of the utterance and therefore changing the thought it expresses*. One of the first philosophers to have made this point was Arthur Prior, who used an example very similar to Katz's 'thank God' example, in support of the opposite conclusion. I already quoted the example a couple of times:

> One says, e.g. 'Thank goodness that's over', and not only is this, when said, quite clear without any date appended, but *it says something which it is impossible that any use of a tenseless copula with a date should convey*. (Prior 1959/1976: 84; emphasis mine)

The same point was to be made forcefully by Castañeda some years later, and, following Castañeda, by John Perry in a sequence of insightful and influential

papers.[84] As a result of their work, the eternalization thesis is as commonly rejected nowadays as it was accepted in the first half of the twentieth century. The irreducibility and indispensability of indexicals is widely acknowledged.

The problem and its solutions

'The essential indexical', Perry says in his well-known essay by the same name,

> is a problem for the view that belief is a relation between subjects and propositions conceived as bearers of truth and falsity. (Perry 1979/1993b : 34)

If we individuate propositions in truth-conditional terms (in such a way that two utterances express the same proposition iff they are true in the same circumstances), then, indeed, the essential indexical poses a problem for the view that belief is a relation to propositions. How can it be that a rational subject believes P while disbelieving Q if P and Q, having the same truth-conditions, are said to be the same proposition ?

To solve that problem there are a number of options available. First, we can make the belief relation triadic : we can say that propositions are believed under 'guises' or 'modes of presentation'.[85] Replacement of an indexical by a non-indexical expression in the asserted sentence affects the guise, even if the proposition expressed is the same. The problem is solved because a rational subject may both believe and disbelieve the same proposition, provided he believes it under one guise (P) and disbelieves it under another guise (Q). This is the solution advocated by Perry himself, as we have seen. He distinguishes two levels of semantic value, corresponding to Kaplan's distinction between character and content, and two relations (belief and acceptance) to those semantic objects. It is therefore possible to

---

[84] See Castañeda's and Perry's respective collections of papers: *The Problem of the Essential Indexical and Other Essays* (Perry 1993b), and *The Phenomeno-Logic of the I: Essays on Self-Consciousness* (Castañeda 1999).

[85] The guise or mode of presentation under which a proposition is believed would, of course, be a function of the guises or modes of presentation of its constituents.

accept P while refusing to accept Q even if P and Q have, or determine, the same content in context.

Alternatively, we can keep the belief relation dyadic, but, departing from Russellianism and the 'coarse-grained' individuation of propositions in terms of objects and properties, follow Frege in building propositions ('thoughts') out of 'senses' or modes of presentation, thus making them directly answerable to cognitive considerations. For that solution to work, special, nondescriptive senses of the sort invoked by the 'neo-Fregeans' must be associated with indexical expressions.[86]

Both of the options I have listed — the neo-Russellian option and the neo-Fregean option — appeal to 'modes of presentation'. That makes them rather similar ; so similar that it may be a mistake to oppose them. Indeed, there are positions that are intermediate between them and cannot be easily classified as either neo-Russellian or neo-Fregean. In earlier work (Recanati 1993, 1995b) I described one such position : it keeps the belief relation dyadic by incorporating the modes of presentation into the singular proposition, alongside the objects and properties of which they are modes of presentation. The resulting 'quasi-singular proposition' is truth-conditionally equivalent to, but cognitively distinct from, the original singular proposition.

Besides this family of view, there is another type of approach, which does *not* posit 'modes of presentation'. That is the sort of view advocated by David Lewis in his treatment of belief *de se*. Lewis's theory of *de se* thoughts can be seen as a generalization of Prior's treatment of tensed sentences. For Prior, as we have seen, tensed sentences express propositions which are true or false only relative to a time. Such propositions are incomplete, by Fregean standards : they are best thought of as propositional functions (taking times as arguments) or as predicates (of times). Incomplete though they are, we can maintain that *they are* the contents of tensed sentences, and also the content of the attitudes (e.g. relief) that we express by using such sentences. The relevant time, without which no truth-value can be determined, is arguably not a part of the content of the sentence, but an aspect of the *circumstance* in which the content is evaluated. We can treat indexical sentences in the same way, by holding that they express relativized propositions : propositions

---

[86] Such senses are of 'limited accessibility' since they « can only be expressed in special circumstances » (Perry 1979/1993b: 45.)

true at some indices but not at others. Thus if John is hungry at t_1, 'I am hungry' is true at <John, t_1>. The proposition expressed by that sentence is a relativized proposition, i.e. a function from centered worlds to truth-values — a property of individuals. Such a proposition is very different from the unrelativized proposition that John is hungry at t_1, hence it is no mystery that one can believe the relativized proposition expressed by 'I am hungry' while disbelieving the unrelativized proposition expressed by 'John is hungry at t_1'. To believe the relativized proposition that one is hungry is, in effect, to self-ascribe the property of being hungry. As we have seen, Lewis generalizes this point and argues that the object of the attitudes, in general, are not (classical) propositions, but properties. Believing a classical, unrelativized proposition P is nothing but to self-ascribe the property of inhabiting a possible world in which it is the case that P.

Even though, in 'Thought without representation' (1986b), he himself appealed to relativized propositions, Perry devoted a section of his earlier paper 'The problem of the essential indexical' (1979) to argue *against* that solution to the problem of the essential indexical :

> The problem [i.e. the problem which the essential indexical raises for the view that belief is a relation to propositions individuated in truth-conditional terms] is not solved... by moving to a notion of proposition that, rather than true or false absolutely is only true or false at an index or in a context (at a time, for a speaker, say). (Perry 1979/1993b : 34)

I will present and discuss Perry's argument in detail in Chapter 35. As I have said already, there is a close connection between that argument and some of the issues raised in the previous chapters, especially the Externality Principle (Chapter 31), and the Reflexive Constraint (Chapter 29).

Perry's argument

Can we use the relativized-propositions framework to deal with the essential indexical ? Can we say that an utterance such as 'I am hungry' expresses an Austinian proposition, the right-hand-side of which is occupied by a relativized proposition, true only at a time and an agent ? According to Perry, if we say so, that will not help us solve the problem of the essential indexical. We cannot, in this way, properly capture the *de se* belief which is expressed by saying 'I am hungry'.

Perry's argument proceeds in two steps. First, Perry attempts to establish that the subject does not merely believe the relativized proposition. The belief could not be evaluated as true or false if its content was exhausted by that proposition. For the belief to be evaluable, we need a situation of concern over and above the relativized proposition. In particular, we need a time and an agent, such that the relativized proposition is *believed to be true with respect to that time and to that agent*. Second step : Perry shows that, as soon as we bring the agent into the picture, the problem of the essential indexical re-appears :

> Once we have adopted these new-fangled propositions, which are only true at times for persons, we have to admit also that we believe them as true for persons at times, and not absolutely. And then our problem returns. (Perry 1979/1993b : 44)

The problem returns because there are different ways of thinking of the person relative to which the relativized proposition is believed to be true. When Perry thinks 'I am making a mess' at time t_1, he believes the relativized proposition 'x is making a mess at t' to be true for himself at that time. But all the shoppers who watch him make a mess also believe that relativized proposition to be true for Perry at t_1. Appealing to Austinian propositions consisting of a situation and a relativized

proposition does not therefore solve the problem. Whether we use classical propositions or Austinian propositions, it seems that we need guises over and above the usual propositional constituents, in order to distinguish Perry's first person belief from the other shoppers' third person beliefs.

Feldman's response

Evidently, one should block the argument at step 1 and maintain that the content of the belief *is* the relativized proposition. Only in that way can we hope to solve the problem of the essential indexical, for the very reason that Perry gives at step 2. The position I take is therefore the same as that defended a long time ago by Richard Feldman in his reply to Perry:

> Perry takes the doctrine of indexed propositions to entail that we do not simply believe such propositions, but rather believe that they are true at some index. However, we need not understand the doctrine in that way. Contrary to what Perrys says, we ordinarily do *not* believe that indexed propositions are true at some index. We simply believe them. In Perry's example, when I realized what was happening I first came to believe the proposition *that I am making a mess*. Of course, I was then believing it a a certain place and time and in a certain possible world. And it was true for me in that world at that place and time. Prior to my realization, I did not believe this proposition at all, although I may have believed some other proposition about this proposition. That is, I may have believed the meta-proposition that the proposition *that I am making a mess* is true at some index, namely, one containing the guilty shopper and that time. Similarly, the shopper watching me does not believe the proposition *that I am making a mess*, but he may believe some proposition about this proposition. For example, he may believe that it is true at an index containing me and then. So, on this view, the proposition *that I am making a mess* is one that I came to believe at the appropriate time and my coming to believe it can help to explain why I straightened my sack. (Feldman 1980 : 82)

Indeed, in the situation-theoretic framework developed by Barwise (1989a), the cognitive content of the belief (that which accounts for the subject's behaviour) is

captured by the right-hand-side in the Austinian proposition, that is, by the relativized proposition. The situation is needed only to account for the belief's truth-conditions. So what the guilty shopper believes, on the situation-theoretic account, is the relativized proposition true at an agent *x* and a time *t* iff *x* is making a mess at *t*.

To be sure, that proposition is not semantically complete : it can be truth-evaluated only with respect to an agent and a time. In his reply Feldman says that the agent and the time are, simply, the agent and time of the context. The agent is the person who believes the relativized proposition, and the time is the time at which the agent believes it. Now Perry had anticipated such a position, and he responded to it in advance :

> All believing is done by persons at times, or so we may suppose. But the time of belief and the person doing the believing cannot be generally identified with the person and time relative to which the proposition believed is held true. (Perry 1979/1993b : 44)

This is the critical issue indeed. According to Feldman, an agent-relative proposition can only be evaluated with respect to the agent in the context of belief, i.e. with respect to the believer himself. So the agent does not have to be represented in order to play its role in fixing the belief's truth-conditions : It is provided by the environmnent. This is also the position defended by Loar, and by Lewis following him. According to Loar, there is a primitive relation, the 'self-ascriptive belief relation', between believers and propositional functions (Loar 1976 : 358). Whenever a person stands in that relation to a propositional function, she entertains a *de se* belief, true iff the propositional function is true-of *that person*. Here again, we find that *the index with respect to which the relativized proposition is evaluated is bound to be the 'index of the context'* (to use the terminology from Lewis 1980). Only if we accept this constraint can we hope to solve the problem of the essential indexical by appealing to relativized propositions.

The (Generalized) Reflexive Constraint

The Constraint I have just stated is a generalization of the Reflexive Constraint from Chapter 29, so I will call it the 'Generalized Reflexive Constraint' (GRC). The

Reflexive Constraint from Chapter 29 said that the proposition that is the content of the state holds in a situation involving the subject of the state — a situation in which the subject stands in a certain relation *R*, fixed by the mode, to what is represented. The Generalized Reflexive Constraint says that the situation with respect to which the *lekton* is evaluated (the Lewisian index) is the situation in which the state is tokened (the context). Not only the subject, but the other coordinates of the index (time, place, etc.) are reflexively identified to the corresponding parameters of the context.

According to Perry, there is no reason to accept the Generalized Reflexive Constraint, for relativized propositions can be held true with respect to indices distinct from the index of the context.[87] That is what Perry thinks happens in the supermarket example : the guilty shopper takes the relativized proposition '*x* is making a mess' to be true with respect to himself, but the other shoppers who watch him also take the relativized proposition to be true with respect to him. The difference — or more cautiously : one difference — between the guilty shopper who holds a first person belief and the other shoppers who hold a third person belief is that, for him but not for them, the 'context of evaluation' and the 'context of belief' coincide. Now the simple fact that they need not coincide shows that the problem of the essential indexical (i.e. the problem of characterizing the first person perspective) cannot be solved simply by appealing to relativized propositions. So the argument goes. As Perry puts it,

> The time of belief and the person doing the believing cannot be generally identified with the person and time relative to which the proposition believed is held true. You now believe that *that I am making a mess* was true for me, then, but you certainly do not believe it is true for you now, unless you are reading this in a supermarket. Let us call *you* and *now* the context of belief, and *me* and *then* the context of evaluation. The context of belief may be the same as the context of evaluation, but need not be. (Perry 1979/1993b : 44)

---

[87] This is Lewis's terminology. Perry would say : the 'context of evaluation' may be distinct from the context of the token state whose content is evaluated.

However, we cannot consider Perry to have demonstrated the possibility of a divergence between the circumstance of evaluation and the context of belief, by providing a couple of examples; for the examples he provides are controversial (to say the least). As we have seen, Feldman *denies* that the other shoppers are belief-related to the relativized proposition '*x* is making a mess at *t*' : they are, at best, related to a meta-proposition about it. Loar and Lewis would make the same denial. So the question we must ask is : Are there good theoretical reasons for accepting, or for rejecting, the Generalized Reflexive Constraint ?

The Externality Principle and the GRC

If the situation which a representation concerns was always fixed by environmental facts like the time of thinking/speaking or the location or identity of the thinker/speaker, that would be sufficient to justify the Generalized Reflexive Constraint. There would be no divergence between the context of belief and the situation of evaluation; the index with respect to which a representation is evaluated would always be the index of the context. But I argued that the situation of concern may be fixed by cognitive factors. One may entertain the place-relative representation 'It is raining' in the course of thinking about a place distinct from the place where one is. In such a case the index relative to which the representation 'It is raining' is evaluated is not the index of the context, because the 'place'-coordinate of the index has been shifted to the place currently under focus.

The Externality Principle and the Generalized Reflexive Constraint turn out to be two sides of the same coin. The Externality Principle tells us that the situation of concern is fixed by environmental facts, not cognitive factors. The Generalized Reflexive Constraint tells us that the situation of concern is anchored to the context and cannot be shifted. Were the Externality Principle correct, it would follow that the Generalized Reflexive Constraint holds and that the sort of divergence between the context of evaluation and the context of belief which Perry invokes in 'the Problem of the Essential Indexical' cannot arise. His argument against relativized propositions would collapse.[88]

---

[88] Here we spot a significant tension in Perry's position (as I have reconstructed it). On the one hand he holds, or seems to hold, that some unarticulated constituents do not belong to the

Be that as it may, I reject the Externality Principle : the utterance/thought 'It is raining' *may* concern all sorts of place, whether or not the speaker/thinker happens to be in that place ; hence there may well be a divergence between the context of belief and the situation of evaluation, as Perry claims. Since I reject the Externality Principle, and the Generalized Reflexive Constraint that goes with it, it seems that I should accept Perry's conclusion : that the problem of the essential indexical cannot be solved by appealing to relativized propositions. Those, like Loar and Lewis, who believe that the problem *can* be solved in this way take the Generalized Reflexive Constraint for granted: what determines the individual with respect to which the self-ascribed propositional function (or property) is evaluated is an environmental fact : the individual in question is bound to be the person who does the self-ascribing (the 'agent' in the 'context of belief'). There is no way in which one can, as it were, vary the person of evaluation by applying the propositional function to someone else. Thus the type of case imagined by Perry — the other shopper's applying the propositional function '*x* is making a mess' to Perry — cannot arise. But for me, given the framework I adopt, such a situation ought to be possible. I may comment on someone's appearance and say : 'Very handsome !' Here, arguably, I express a person-relative proposition (a property), true of persons at times, and I apply it to a certain individual whom my utterance concerns. Whether or not such a view is sustainable, it is clearly in the spirit of the general position I have put forward. Now if we accept that *there are* such person-relative propositions, which can be evaluated with respect to whichever persons they happen to concern, then it is clear that the problem of the essential indexical cannot be solved merely by appealing to such propositions. For the person whose appearance I comment upon when I say (or think) 'Very handsome !' may happen to be myself, seen in a mirror and mistaken for someone else. In such a situation, arguably, I believe the propositional function '*x* is very handsome' of myself, yet I do not believe the sort of thing that I could express by

---

situation of concern because they are not fixed by environmental factors but by cognitive factors. He therefore accepts, or seems to accept, the Externality Principle. On the other hand, when he insists that the problem of the essential indexical cannot be solved by appealing to relativized propositions, he argues that the Generalized Reflexive Constraint must be rejected. There is a tension because, as I have said, the Externality Principle and the GRC are two sides of the same coin.

saying 'I am very handsome'. I entertain a *de re* belief about myself, not a *de se* belief. I conclude that, without something like the Generalized Reflexive Constraint to anchor the situation of concern to the context, the problem of the essential indexical cannot be solved by appealing to relativized propositions (unsupplemented by guises or something of that sort).

In Chapter 36, this objection will be met, by accepting that — as we have already seen (Chapter 29) — the Generalized Reflexive Constraint holds in a limited range of cases.

Indexicality and the GRC

One of the differences between 'It is raining' and 'It is raining here' is that, while the first sentence *can* be evaluated with respect to a place different from the place of the context, the second sentence cannot : as we have seen (chapter 12), the indexical 'here' rigidly anchors the situation of evaluation to the context.[89] Indexicals are freezing operators that make the relevant aspect of the situation of evaluation unshiftable. As far as indexicals are concerned, therefore, the Generalized Reflexive Constraint holds : the coordinates of the evaluation index that correspond to indexicals cannot be shifted but are set, once for all, by the context.

It is interesting to note that Feldman, who maintains the Generalized Reflexive Constraint against Perry, models index-relative propositions by indexical sentences. To believe a relativized proposition, for him, is to accept an indexical sentence such as 'I am making a mess' :

> The idea [of believing a relativized proposition] may be clarified if we drop talk of propositions altogether and simply talk of sentences. What I came to believe was the sentence 'I am making a mess'. I did not believe it before. I came to believe it at the same time I became prepared to say it. Note that I need not have become prepared to say-it-at-an-index or say that it is true at an index. I may have no thoughts about indices at all. I just became prepared to say this sentence. Of course, there was a time, place, etc., at which this happened. Similarly, I did not come to believe that the sentence is true at an index. I simply came to believe it. This coming to believe may have occurred at some index, but the index is not in any sense a part of the content of my belief. (Feldman 1980 : 82-3)

---

[89] There are apparent exceptions to that principle (e.g. shifted indexicals in free indirect speech) but I will not deal with them here. See Recanati forthcoming b.

Since he equates relativized propositions with indexical sentences, it is understandable that Feldman sticks to the GRC. For it is a property of indexical sentences that the relevant coordinates of the index are anchored to the context and cannot shift.

At this point someone like Feldman, who thinks the problem of the essential indexical can be solved by appealing to relativized propositions, may argue as follows. It is true that the unarticulated constituents which go into the evaluation index are shiftable ; hence if we think of *de se* belief on the model of 'It is raining' the problem of the essential indexical will not be solved, because the GRC will not hold. But we need not think of *de se* belief on the model of 'It is raining'. We may think of it on the model of 'It is raining here'. This is an indexical sentence, and that guarantees that the GRC holds.

This intriguing move raises an obvious objection. In contrast to 'It is raining', the sentence 'It is raining here' does *not* express a place-relative proposition. Being articulated (by the indexical 'here'), the place of rain is a constituent of the content articulated by the sentence — it is a constituent of the *lekton*. To say that *de se* belief should be thought of on the model of 'It is raining here' rather than on the model of 'It is raining' is therefore to *give up* the view that *de se* belief should be accounted for in terms of relativized propositions.

Obvious though it is, the objection can be met, by distinguishing two dimensions under which indexical sentences differ from sentences like 'It is raining'. In terms of this distinction, the proposal can be rephrased : the suggestion is that *de se* thoughts are more like indexical sentences than like sentences such as 'It is raining' *as far as one of the two dimensions is concerned*. The situation is reversed when one considers the other dimension.

(Un)articulatedness and (un)shiftability

Indexical sentences such as 'It is raining here' contrast with location-relative sentences such as 'It is raining' in two dimensions. First, the place is articulated in the indexical sentence 'It is raining here', while it is unarticulated in 'It is raining'. In my framework, whatever is unarticulated falls on the situational side, hence 'It is raining' expresses a location-relative proposition — a proposition true at some

locations and false at others. Second, the place which the indexical adverb 'here' articulates cannot be shifted ; it is bound to be the place of the context. This is a property of the indexical qua freezing operator. In contrast, the place which remains unarticulated in 'It is raining' and features in the situation of evaluation can be shifted : as we have seen, it need not be the place of the context but can be any place that is currently in focus.

The suggestion, then, can be rephrased as follows. Let us construe the two dimensions as independent. They determine a matrix with four cells (Table 2). On the articulated side, we find not only the indexicals, which have the property of unshiftability, but also pronominal expressions which do not have that property.[90] On the unarticulated side, a sentence like 'It is raining' expresses a location-relative proposition that can be evaluated with respect to any place, as we have seen ; but why should there not also exist sentences which, like 'It is raining', express location-relative propositions, but propositions which can only be evaluated with respect to the place of the context ? Why should the fourth cell in the matrix (unarticulatedness + unshiftability) remain empty ? What I suggest, therefore, is that we make room for a new category : that of *context-relative representations*, representations which display the two features of relativity (unarticulatedness) and unshiftability. Relativity makes them similar to 'It is raining', and unshiftability to 'It is raining here'.

	shiftable	unshiftable
Articulated	*pronouns*	*indexicals*
Unarticulated	*situation-relativity*	*context-relativity*

Table 2

If the fourth cell is not empty, there is a class of representations for which the Generalized Reflexive Constraint holds, even though they are not indexical.

---

[90] 'It is raining here' can only refer to the place of the context ; but if we substitute the (arguably) pronominal adverb 'there' for the true indexical 'here', the situation changes : 'it is raining there' can refer to *any* place, and can even be used in quantificational environments, where no specific place is at issue ('Whenever I arrive in a faraway place after a long and tiresome trip, it is raining there').

I will henceforth use 'context-quotes' to conventionally indicate that a given representation is relative to some feature of the context in which it is tokened. Thus in contrast to 'It is raining' which, because it can be evaluated with respect to any place, is relative to a *shiftable* situational component, ^{c}It is raining^c (with the context-quotes) is relative to an *unshiftable* situational component : it can only be evaluated with respect to the place of the context. That is a feature it shares with the indexical sentence 'It is raining here'. But in contrast to 'It is raining here', the relevant situational component (that to which the proposition is said to be relative) is not articulated, indexically or otherwise, in ^{c}It is raining^c.

Context-relative representations in language and thought

Context-relative representations, if they exist, are *intermediate* between indexical representations and situation-relative representations such as 'It is raining' (which can be evaluated with respect to any place and time). Like indexical sentences, their truth-value depends upon a feature of the context — for example the time or the place of the tokening. Whether we evaluate the indexical sentence 'It is raining here and now', or the context-relative sentence ^{c}It is raining^c, in both cases we must look at the place and time of the context to check whether or not the sentence is true. But the relevant feature of context is explicitly represented in the indexical case, while it remains unarticulated in the context-relative case, just as the situation of concern remains unarticulated in situation-relative sentences.

Do such representations exist? As far as natural language is concerned, that is far from obvious. It may be that, in natural language, we find only indexical sentences ('It is raining here and now') and situation-relative sentences ('It is raining'), but no context-relative sentences. Be that as it may, we have already seen that *there are* context-relative representations in thought. The most basic kind of representation with which a perceiving-and-acting organism must be credited presumably belongs to that category.

An organism which (like most animals) does not have the reflective capacity to think of itself as a person among other persons yet perceives and acts should be credited with a primitive form of egocentric thinking — what I called a 'basic *de se*

thought' in Chapter 24.[91] That is one of Perry's major insights, in 'Thought without Representation' and elsewhere. I already quoted the following passage from 'Perception, Action, and the Structure of Believing':

> The information that we get at a certain spot in the world is information about objects in the neighborhood of that spot in a form suitable for the person in that spot. As long as this is the only source of information we have about ourselves, we need no way of designating ourselves, indexical or insensitive. Our entire perceptual and doxastic structure provides us with a way of believing about ourselves, without any expression for ourselves. (Perry 1986a/1993b : 148-149)

Though relative to the subject, perceptual representations are not just person-relative representations ; for they are not applicable to other persons than the subject of the context. They can *only* be evaluated with respect to the context in which they are entertained — the perceiving subject, and the time and place of perception. In other words, they are context-relative representations, satisfying the Generalized Reflexive Constraint.[92]

Shall we say that only organisms devoid of an explicit representation of themselves entertain such context-relative representations, while we entertain more sophisticated, indexical representations containing the word 'I' or a mental analogue? Perry gave us a reason not to make that move. Just as 'there is a little of the Z-lander in the most well-traveled of us', there is a little of the simple perceiving-and-acting organism in the most reflectively self-conscious of us. What Perry says about our

---

[91] Millikan (2004 : 179) calls this an 'ego-implicit' representation.

[92] What makes them context-relative is a feature distinct from their content, namely their mode. (The perceptual mode of representation is what forces the index of evaluation to be the index of the context, in conformity to the GRC.) Hence it was a mistake to talk of 'context-relative *propositions'*, as I did in 'Relativized Propositions' (Recanati 2006b): only *representations,* endowed with a mode in addition to a propositional content, can be said to be context-relative.

thoughts concerning the local weather easily generalizes to perception-based thought, whether it is about the weather or anything else:

> What each of us gets from perception may be regarded as information concerning ourselves, to explain connections between perception and action. There is no need for a self-referring component of our belief, no need for an idea or representation of ourselves. When a ball comes at me, I duck ; when a milk shake is put in front of me, I advance. The eyes that see and the torso or legs that move are parts of the same more or less integrated body. And this fact, external to the belief, supplies the needed coordination. The belief need only have the burden of registering differences in my environment, and not the burden of identifying the person about whose relation to the environment perception gives information with the person whose action it guides. (Perry 1986b/1993b : 219)

Again, perceptual representations are context-relative. They satisfy the GRC, but they do so in virtue of architectural facts, not through the use of explicit indexicals.

Expressing one's subjective thoughts

If there are context-relative representations in thought, but (presumably) not in language, how can we verbally express a context-relative thought, e.g. the thought that one is hungry or cold, or the thought that the wine tastes bad ?

If we want to express a context-relative thought, using language, we have to use either a situation-relative sentence or an indexical sentence. If a situation-relative sentence is available that is the most appropriate tool. The difference between a situation-relative sentence and a context-relative sentence is small when the situation of evaluation happens to be the situation of utterance. It is a purely modal difference : the situation-relative sentence *could be* used to characterize another situation than the situation of utterance, while a context-relative sentence could not. The difference between a context-relative sentence and an indexical sentence is more substantial,

since the indexical sentence articulates the relevant contextual feature and makes it part of the *lekton*. That is why one would not express the context-relative thought ^cIt is raining^c, prompted by the perception of rain, by saying 'It is raining here and now'. We would naturally say 'It is raining' (unless there is a good reason to explicitly represent the time and place of the context).

But there are context-relative thoughts for which no situation-relative sentence is available. Self-relative thoughts are a case in point. If I want to express a self-relative thought, to the effect that I am hungry or cold, I can hardly say 'Hungry !' or 'Cold !', for that is not proper English. I have to use an indexical sentence and say : 'I am cold', or 'I am hungry'. In contrast to 'It's raining here', which makes sense only if there is a good reason to explicitly represent the place of the context (typically because of an intended contrast between that place and some other place), I can use an indexical sentence such as 'I am hungry' to express a context-relative thought because, given the lack of an appropriate situation-relative sentence, there is no clear alternative. In this way we can perhaps explain why Wittgenstein and some of his followers have insisted that there are two uses of 'I' : a subjective use whereby an indexical sentence 'I am F' expresses a self-relative thought, and a more objective use whereby it expresses a proposition with the subject as a constituent. The word 'I', they say, is a genuine referring expression only in the second type of case (see e.g. Chauvier 2001).

Sometimes, the language makes a situation-relative sentence available to express a self-relative thought. 'The wine tastes bad' is an example. I can use that sentence to express the self-relative thought ^cThe wine tastes bad^c. I do not have to say : 'The wine tastes bad to me'. Still, the thought I express *is* irreducibly subjective : ^cThe wine tastes bad^c — the thought I express — is true if and only if the wine tastes bad *to me*, the thinker of that thought. In contrast to the thought, however, the natural language sentence 'The wine tastes bad' is not itself context-relative, but only situation-relative. As Kölbel and Lasersohn have argued, predicates of personal taste require a 'judge' coordinate in the index of evaluation, but the judge need not be the speaker — it may be any person or group whose taste serves as standard.[93] So the

---

[93] « In certain circumstances, we may… adopt an *exocentric* perspective, assessing sentences for truth relative to contexts in which someone other than ourselves is specified as the

claim I make when I say 'The wine tastes bad' can always be interpreted as more general or more objective than it is (i.e., than it is when construed as the mere expression of a self-relative thought). If I want to avoid such interpretations of my claim, I will have to make the relativity to myself explicit by using an indexical : 'The wine tastes bad to me'. But if I do that, I introduce a possibly unwelcome contrast between my viewpoint (my taste) and that of others, a contrast which is absent when I merely express my feeling.

---

judge… For example, suppose John is describing to Mary how their two-year old son Bill enjoyed a recent trip to the amusement park. Something like the following dialog might occur :

> Mary : How did Bill like the rides ?
> John : Well, the merry-go-round was fun, but the water slide was a little too scary.

We intuitively regard John's utterance as true if the merry-go-round was fun for Bill, not if it was fun for ourselves (or for John). » (Lasersohn 2005 : 670-72)

Chapter 37

Basic and nonbasic *de se* thoughts

Return of the essential indexical

I have just mentioned the Wittgensteinian distinction between 'two uses of 'I''.'I' is used 'as subject' in a sentence 'I am F' if the sentence expresses a self-relative thought, as when I feel hot and I express my feeling by saying 'I am hot'. It is used 'as object' if the sentence expresses a thought explicitly about myself, e.g., the thought that I was born in 1952. In my framework, a thought explicitly about oneself is a thought the content of which (the *lekton*) involves oneself as a constituent. That is not sufficient, of course : such a thought counts as a first person thought, as a thought about *oneself*, only if the constituent in question is thought under the mode of presentation **EGO**.

This connects up with the distinction I drew earlier between basic and nonbasic *de se* thoughts (Chapter 24). In basic *de se* thoughts, the self is not articulated, it only features in the index of evaluation. Basic *de se* thoughts are, in effect, self-relative thoughts. In nonbasic *de se* thoughts, the self is explicitly identified, and it is identified *as oneself*. Nonbasic *de se* thoughts contrast both with basic *de se* thoughts, which are identification-free (they 'concern' the subject without being about him or her), and with those *de re* thoughts which happen to be about oneself but which are not internally first personal (because one is not thought of under the mode of presentation **EGO**, but under a third person mode of presentation, as in Kaplan's mirror example).

The distinction between the two uses of 'I', and the two types of *de se* thought, leads us back to the problem of the essential indexical. I have shown that, by appealing to context-relative representations, we can account for *de se* belief. In defense of a Lewisian type of account, I have argued that Perry's objection does not hold, since context-relative representations satisfy the Generalized Reflexive Constraint. But the problem of the essential indexical is not thereby solved in its full

generality, for it arises also with respect to the 'objective' use of indexical sentences. If I entertain an explicitly indexical (rather than context-relative) thought and express it by uttering the sentence '*I* was born in Paris (in contrast to you, who were born in Chicago)', my utterance expresses a proposition with me as constituent. Yet the indexical is no less essential in that sentence than in any other : if we replace it by a non-indexical expression, we affect the cognitive significance of the utterance. As Perry might say, 'our problem returns'.

To sum up, the Lewisian account with its relativized propositions is appropriate only for basic *de se* thoughts which are, indeed, self-relative. But basic *de se* thoughts are only one particular category of *de se* thoughts. Another category is that of nonbasic *de se* thoughts, and for them, an account à la Perry, with first person modes of presentation, seems unavoidable. No 'EGO' mode of presentation is required to account for basic *de se* thoughts, which are identification-free and involve no explicit representation of the self ; but we need such a mode of presentation if we are to come up with a *general* account of *de se* thought, and a general solution to the problem of the essential indexical.

De re *thoughts as a variety of* de se *thoughts*

In 'Attitudes de dicto and *de se*' (section XII), Lewis has compared his own account of indexical belief with that of Perry, and he puts the matter thus : For Perry, *de se* thoughts are a species of *de re* thoughts ; for Lewis, it is the other way round. Lewis summarizes Perry's position as follows:

> Perry's scheme (…) provides, in the most straightforward way possible, for other-ascription as well as self-ascription of properties. Ascription of properties to individuals, in general, is called belief *de re*. Perry's scheme is made for belief *de re,* and belief *de se* falls under that as a special case. (Lewis 1979/1983 : 151-2)

For Lewis, however, « The subject's self-ascriptions are the whole of his system of beliefs. Other-ascriptions of properties are not some further beliefs alongside the self-ascriptions », but a particular case of self-ascription.

Let us start with *de re* thoughts — thoughts about a certain object *x*, to the effect that it has property *F*. Such thoughts give rise to *Frege's problem* : it is possible for a rational person to think of the same object, *x*, both that it is *F* and that it is not *F*. How can that be ? As Frege pointed out, the problem is solved by appealing to modes of presentation over and above the objects thought about. A rational subject can believe of *x*, thought of under a mode of presentation *m*, that it is *F*, and at the same time believe of the same object *x*, thought of under a different mode of presentation *m'*, that it is not *F*. Insofar as the modes of presentation are distinct, there is no irrationality. So *de re* thoughts involve two things : the *res* thought about, and the way it is thought about (the mode of presentation). The mode of presentation is the internal, psychological aspect of the thought, while the *res* thought about depends upon the external environment.

De re thoughts in general involve *nondescriptive* modes of presentation, that is, modes of presentation which essentially depend upon contextual relations to the object. The object the thought is about is the object which stands in the right contextual relations to the subject or to the mental episode in which the mode of presentation occurs. For example, a demonstrative mode of presentation 'that *G*' is typically based upon a certain contextual relation to the object, in virtue of which the subject is able to focus his or her attention on it and gather (typically perceptual) information from it. The mode of presentation can be thought of as a temporary mental file in which the subject can store the information acquired in virtue of the contextual relation in question.[94]

According to Perry, the concept **EGO** is a special case. It is a mental file that is based upon a special relation which every individual bears to himself or herself, namely identity. In virtue of *being* a certain individual, I am in a position to gain information concerning that individual in all sorts of ways in which I can gain information about no one else, e.g. through proprioception and kinaesthesis.[95] The mental file **EGO** serves as repository for information gained in this way (as well as for any information that is about the same individual as information gained in this way).

---

[94] See Recanati 1993, 2006a for a more detailed presentation of this familiar account.

[95] As Frege says, « every one is presented to himself in a particular and primitive way, in which he is presented to no one else » (Frege 1918-19: 25-6).

So *de se* thoughts are a variety of *de re* thoughts, namely *de re* thoughts involving the mode of presentation ᴇɢᴏ.

Lewis disagrees. For Lewis, *de se* thoughts are not a species of *de re* thoughts, but the other way round. Let us assume, following Perry (and Evans), that a nondescriptive mode of presentation of the sort that occurs in *de re* thoughts is constitutively associated with a certain relation R to objects, a relation whose obtaining creates an information link that the subject can exploit. This story is compatible with Lewis's claim that *de re* thought is a special case of *de se* thought. When I think a *de re* thought to the effect that a certain object is F, and the object in question is thought of under a mode of presentation m based upon a certain relation R_m to that object, what I do, according to Lewis, is *self-ascribe the property of standing in relation R_m to an x that is F*. For example, to think 'that is F' is to self-ascribe the property of standing in the demonstrative relation R_{dem} to some object x that is F (Lewis 1979/1983 : 154-55). By contrast, when I think a *de se* thought such as 'I am F', I directly self-ascribe the property of being F. On this account, *de re* thoughts are a variety of *de se* thought, but not the other way round.

<p align="center">De se thoughts as a variety of de re thoughts</p>

I see two problems with Lewis's view. First, the view entails, or seems to entail, that the subject is able to think of the contextual relation R in which he or she stands to the object of his/her thought. This entailment is absent from the pure 'mode of presentation' view. That view simply says that the subject thinking a *de re* thought exercises a mental file, whose existence is contingent upon a certain relation R to the object. The subject need not be aware of that relation, or be able to articulate it. So for example the relation can be that which holds when a Kripkean causal chain links me to Cicero. We may be reluctant to say that the subject who believes that Cicero was rich thereby self-ascribes the property of being related to some rich individual x by a K-causal chain involving his use of the name 'Cicero', since the subject presumably lacks the notion of a K-causal chain.

Let us, however, grant that this objection can be met. The major problem I see with Lewis's view is that he fails to draw the crucial distinction between basic and nonbasic *de se* thought. In terms of that distinction, it is clear that Lewis's point is that *de re* thought is a species of *basic de se* thought. The self's involvement is implicit : it

is determined by the belief mode (construed as a self-ascriptive mode) and is invariant. The invariance in question is of the strongest form possible: every belief, for Lewis, is *de se*. Whether the subject thinks 'I am hot', 'That is gigantic', or 'Ice melts', in all cases the subject self-ascribes a property : the property of being hot, the property of being R_{dem}-related to an object x that is gigantic, or the property of inhabiting a possible world in which ice melts.

What about nonbasic *de se* thought, i.e the case in which the subject explicitly thinks about himself/herself ? What about a thought like 'I was born in 1952' ? It is here that Lewis's view and Perry's seem to conflict. Lewis offers his theory as a theory of *de se* thought *in general.* For him, the subject who thinks 'I am Hume' self-ascribes the property of being Hume. Presumably, he would hold that the subject who thinks 'I was born in 1952' self-ascribes the property of having been born in 1952. This, for Lewis, is *not* a *de re* thought. Indeed, no Frege case can arise when I self-ascribe a property : I cannot be misled as to whom I ascribe the property to, because the identity of the individual to whom I ascribe the property is determined by the architecture of the belief system in an invariant manner. A belief is a *self*-ascription, period. To use Millikan's terms, the ascribee is simply not among the 'variables that can accept alternative values'. When I hold an ordinary belief, the self is represented only implicitly, just as the actual world is represented only implicitly.

Lewis does not distinguish nonbasic from basic *de se* thought, and that is the main shortcoming of his theory. For there *are de se* thoughts in which the self is represented explicitly. (They are what Millikan 2004 calls 'ego-explicit' thoughts.) Such thoughts naturally give rise to Frege cases. Imagine that I see myself in the mirror and come to believe that I have a bump on my forehead. Two modes of presentation are involved here : a demonstrative mode of presentation of the man I see in the mirror, and the **EGO** mode of presentation. I am related to the individual my thought is about both through the R_{dem} relation and the R_{ego} relation. The information that I have a bump on my forehead is gained through the R_{dem} relation (it is gained by looking at the mirror), but I know that the person I am looking at is myself, so I safely transfer that information into my **EGO** file, and can act in the relevant way (by touching my forehead and checking the presence of the bump). Since modes of presentation are involved, mistakes can be made. I can be wrong in assuming that the man in the mirror is myself ; or, on the contrary, I may fail to recognize the man in the mirror as myself and entertain contradictory beliefs about what is in fact the same person

('That man has a bump on his forehead, but I do not'). All this shows that the *de se* thought 'I have a bump on my forehead', in this type of case, is a species of *de re* thought, as Perry claims. But this is compatible with Lewis's view that *de re* thought is a species of *de se* thought. The two views are consistent, given the distinction between basic and nonbasic *de se* thought.

Lewis's claim should be rephrased as the claim that every thought is a basic *de se* thought. The simplest of such thoughts are the self-ascriptions of properties like being hot or being hungry. But more complex properties can be involved, like the property of seeing an object that is *F*, or more generally the property of standing in a contextual relation *R* to some object that is *F*. Thus *de re* thoughts turn out to be a species of basic *de se* thought. Still, we can maintain, with Perry, that nonbasic *de se* thoughts are *themselves* a species of *de re* thought, involving a special relation R_{ego} to the *res* the thought is about. If this is right, then a nonbasic *de se* thought (e.g. the thought that one was born in 1952) is itself a species of basic *de se* thought : to believe that one was born in 1952 is to self-ascribe the property of being R_{ego}-related to an *x* such that *x* was born in 1952.

Since the R_{ego}-relation is the relation of identity, is it easy to overlook the difference between self-ascribing the property of being born in 1952 and self-ascribing the property of being R_{ego}-related to an *x* such that *x* was born in 1952. But there is a significant difference : in one case but not in the other, the object of which *being born in 1952* is predicated is explicitly identified, and can be mis-identified. This is enough to justify talk of 'modes of presentation'.

I conclude that Perry and Lewis are both right. There is a sense in which *de se* thought is a special case of *de re* thought, but there is also a sense in which *de re* thought is a special case of *de se* thought. The apparent conflict between Perry's and Lewis's perspectives is due to insufficient appreciation of the crucial distinction between basic and nonbasic *de se* thoughts. Given that distinction, we need both relativized propositions and modes of presentation in a complete account of belief *de se*.

Part Eleven

Shiftability

Chapter 38

The Generalized Reflexive Constraint

Lewis and the GRC

In Chapter 29 I introduced the Reflexive Constraint as a characteristic feature of perception and memory and, more generally, of a family of states which, following Higginbotham, I called the 'reflexive states' (Chapter 22). In virtue of the Reflexive Constraint, the content of such a state has to be evaluated in a situation involving the subject of the state — a situation in which the subject stands in some relation R, determined by the mode, to what the state represents.

The Generalized Reflexive Constraint (GRC), introduced in Chapter 35, is more demanding. According to the GRC, the situation with respect to which the content of the state is evaluated must involve not only the subject of the state, but also the time of the state, the location of the state, and so on and so forth : in Lewis's terms, the index with respect to which the content is evaluated must be the 'index of the context'. The Generalized Reflexive Constraint anchors the situation of evaluation to the context and makes it unshiftable.

In Lewis's theory, the GRC holds across the board. Whenever a representation is tokened in the subject's mind, the content of the representation is evaluated with respect to a situation involving the subject and the time and place of the tokening. In other words, the content is construed as a property which is ascribed to the context in which the representation is tokened.[96] If the subject feels that she is hot, the content of the representation is the property of being hot, which the subject ascribes to herself at the present time. If the subject believes that this is a hot saucer, the content of the representation is another property which the subject ascribes to herself at the present time, namely the property of being R_{dem}-related to an x such

---

[96] As I pointed out in Chapter 15 (p. 00), the subject who self-ascribes a property is the same thing as a Lewisian context.

that *x* is a hot saucer. If the subject conjectures that it was hot two weeks ago, she conjecturally ascribes to the context she is in the property of being two weeks later than a situation in which it was hot at the same place.

So for Lewis the index of evaluation can only be the index of the context ; it cannot be shifted. To be sure, Lewis acknowledges that the index can be shifted through the use of operators, but, precisely because this is done by means of operators, only sentences *in the scope of operators* can have their index of evaluation shifted. Autonomous sentences are always evaluated with respect to the index of the context. It follows that the GRC holds for them. Embedded sentences are another story : they are evaluated with respect to whichever shifted index may have been introduced by an operator higher up in the (autonomous) sentence of which the embedded sentence is a part.

Operator-shiftability vs *free shiftability*

The claim that the situation of evaluation can be shifted only through the use of operators can be disputed. As I pointed out in Chapter 12, a simple, unembedded sentence like 'It is raining' may well concern a place distinct from the place of utterance. If the speaker and the hearer are involved in a conversation about what is going on at a certain place *l*, distinct from the place of utterance, the sentence 'It is raining' uttered in that context will presumably be about *l*, not about the place of the context. In this regard, 'It is raining' behaves differently than the indexical sentence 'It is raining here', since 'It is raining here' can only be evaluated with respect to the place of utterance. The GRC holds with respect to the latter, it seems, but not with respect to the former. That means that, unless a freezing operator like the indexical adverbs 'here' and 'now' keeps the evaluation relative to the context of utterance, the situation of evaluation is *shiftable* even in the absence of operator.

Lewis accepts only one type of shiftability, however : operator shiftability, i.e. shiftability by means of operators. Modal and temporal operators admittedly shift the point of evaluation : they take us to a point of evaluation distinct from the initial or default point which is the index of the context. But operator shiftability only concerns sentences in the scope of operators ; it does not concern autonomous sentences. The possibility of shifting the index of evaluation for *autonomous* sentences as in the example I have just given is something that Lewis does not acknowledge. Let us call

the form of shiftability which Lewis does not accept, and which a simple sentence like 'It is raining' illustrates, 'free' shiftability.

Faced with prima facie examples of free shiftability such as 'It is raining' used to talk about a remote situation, Lewis can argue that they involve an *implicit* operator. To change the example a bit, suppose that we are talking about the situation in Chicago, and I say : 'At this time of the year, it is freezing'. This may be construed as elliptical for a more complex sentence : 'In Chicago, at this time of the year, it is freezing', where 'in Chicago' is a place-shifting operator taking us from the place of utterance to the place named by the proper noun 'Chicago'.

Lewis himself appeals to implicit operators in his theory of fiction :

> If I say that Holmes liked to show off, you will take it that I have asserted an abbreviated version of the true sentence « In the Sherlock Holmes stories, Holmes liked to show off ». As for the embedded sentence « Holmes liked to show off », taken by itself with the prefixed operator neither explicitly present nor tacitly understood, we may abandon it to the common fate of subject-predicate sentences with denotationless subject tems : automatic falsity or lack of truth value, according to taste. (Lewis 1978/83 : 262-3)

So 'Holmes liked to show off' turns out to be ambiguous, according to Lewis. Taken as an abbreviated sentence headed by an implicit operator, it is true ; taken as unabbreviated, it is (say) truth-valueless. Similarly, on the suggested analysis, 'At this time of the year, it is freezing' is ambiguous. Taken by itself with the prefixed operator 'In Chicago' neither explicitly present nor tacitly understood, it expresses a falsehood if it is uttered in Berkeley in February, because without the operator the sentence is bound to be evaluated with respect to the index of the context (in virtue of the GRC).[97] With the implicit operator prefixed to it, it expresses a truth, because then the embedded sentence in the scope of the operator is evaluated with respect to a shifted index in which Chicago substitutes for Berkeley.

Just as there are autonomous sentences that are, or seem to be, evaluated with respect to a shifted index (e.g. 'It is freezing' used to talk about the situation in Chicago), there are mental representations which do not satisfy the GRC because

---

[97] I assume that it is not freezing in Berkeley in February.

their content is evaluated in a situation distinct from the situation in which they are tokened. Memory states are an obvious example. To be sure, memory satisfies the Reflexive Constraint : the content of a memory must be evaluated with respect to a situation experienced by *the subject who remembers*. But memory does not satisfy the *Generalized* Reflexive Constraint : the situation of evaluation reflexively involves the subject who remembers, but its time coordinate is not the time at which the memory occurs. If, at a time t, I remember that p, it must have been the case that *p at a time t' anterior to t*. Here *t'* is the time of evaluation, and it is distinct from *t*, the time of the memory.

Here also, it is possible to dispose of the apparent counterexample to the GRC, by arguing that the index is shifted through the use of an operator. The memory mode, it may be argued, is in effect an operator analogous to the past tense operator. A memory that p is evaluated with respect to a shifted time *t'* anterior to the time *t* at which the memory occurs. The mode contributes an operator like 'in my past experience, it was the case that'. The content of the sentence embedded under that operator is evaluated with respect to a shifted index, but the content of the complete sentence headed by the operator is evaluated with respect to the index of the context, as the GRC demands.

Farewell to the GRC

Considered as an index-shifting operator, the memory mode M, in a memory representation M(p), contributes a specification of the relation R to which the subject stands to the state of affairs remembered. Instead of remaining external to the content of the representation, the relation R is now internalized and becomes an aspect of the *lekton*, contributed by the operator. That is what makes it possible for the *lekton* to be evaluated with respect to the subject's present situation, in conformity to the GRC, rather than with respect to some past situation in which the subject experienced the scene he now remembers. Only the embedded part of the representation (the proposition p corresponding to what is remembered) is evaluated with respect to the past situation, in the course of evaluating the complete *lekton* (including the M operator) with respect to the present situation.

In chapter 37 I objected to a similar move concerning the nondescriptive modes of presentation that occur in *de re* thoughts. A nondescriptive mode of

presentation is constitutively associated with a certain relation R to objects, a relation whose obtaining creates an information link that the subject can exploit. Here also Lewis's theory *internalizes* the relation R and makes it part of the *lekton*. When I think a *de re* thought to the effect that a certain object is F, and the object in question is thought of under a mode of presentation m based upon a certain relation R_m to that object, what I do, according to Lewis, is *self-ascribe the property of standing in relation R_m to an x that is F*.

I objected to this view that it entails, or seems to entail, that the subject is able to think of the contextual relation R in which he or she stands to the object of his/her thought. This is objectionable because the subject need not be aware of that relation, or be able to articulate it. For example, an ordinary subject who believes that Cicero was a Roman orator presumably lacks the notion of a Kripkean causal chain, so we may be reluctant to say that in so believing he or she self-ascribes the property of being related to some Roman orator x by a Kripkean causal chain involving his use of the name 'Cicero'.

It is unclear whether or not a similar objection to Lewis's story can be raised in the case of memory and other psychological modes which are prima facie counterexamples to the GRC. Be that as it may, I think the original objection — that which concerns modes of presentation — suffices to make us suspicious of Lewis's version of Moderate Relativism. Lewis's theory, characterized by the context-independence of the *lekton* and the GRC, is unattractive because (like Reflexivism) it puts too much into the *lekton*.

On the alternative theory I advocate, the *lekton* is context-dependent, and it is identified to the content in the sense in which 'content' contrasts with 'mode'. Thus the content of a memory that p is simply the proposition that p. That proposition is not evaluated with respect to the subject's present situation, as the GRC demands, but it is evaluated with respect to some situation the subject experienced in the past. On this version of Moderate Relativism, the Generalized Reflexive Constraint does *not* hold, and there is no reason not to acknowledge the phenomenon of free shiftability.

Chapter 39

Parametric invariance and *m*-shiftability

Egocentrism

In Lewis's framework, as we have just seen, every thought is *de se*. Whether the subject thinks 'I am hot' or 'Frege died in 1925', in both cases he or she self-ascribes a property (the property of being hot, or the property of inhabiting a world in which Frege died in 1925). So there is a dissymmetry between 'I am hot' and 'He is hot'. The subject who thinks 'I am hot' self-ascribes the property of being hot. The subject who thinks 'He is hot' self-ascribes the property of being R_{dem}-related to an *x* such that *x* is hot. While Perry treats self-ascriptions and other-ascriptions as being on a par — as being two types of *de re* thought, involving distinct kinds of modes of presentation — Lewis considers the self-ascriptive case as basic, and other-ascriptions as a particular case of self-ascription. Every ascription, including other-ascription, is (at bottom) a self-ascription, for Lewis.[98]

In the Lewisian framework all the thoughts of a given individual concern himself or herself ; hence the invariance condition we talked about in Chapter 33 is satisfied. The self has the same status as Z-land in Perry's tale. Indeed, Lewis's treatment of belief *de se* conforms to Perry's instructions :

> In cases in which *the same unrepresented parameter* is relevant to a whole
> mode of thinking or discourse, we should classify each specific belief or
> utterance with a propositional function. The truth-value would be that of the
> proposition obtained by applying the function to the value of the parameter
> fixed by facts about the whole system. (Perry 1993a: 221 ; emphasis mine)

This invariance claim should be qualified, however. First, the subject to whom the thought is (self-)ascribed is fixed and invariant *only for a given subject (!)*. If we

---

[98] For a comparative discussion of Lewis's and Perry's views of the relations between *de re* and *de se*, see Chapter 37.

consider *several* subjects, then, of course, the subject will vary, and the content of a thought — the property it expresses — will be (self-)ascribed to *the subject of that thought*. Second, even though Lewis speaks of 'self-ascription', what the *lekton* is predicated of in his framework is more accurately described as a *centered world*, involving not only the subject, but a certain world and a certain location in that world. Now the world which the subject inhabits is invariant, in Lewis's system[99] ; but his or her spatio-temporal location in the world constantly varies. Even if we fix the subject and (consequently) the world, the subject's thoughts will concern different centered worlds if they are tokened at different times.

Despite this lack of invariance, something is fixed even in this case : the function that takes the context of thought as argument and returns the situation which the thought concerns as value. In Lewis's framework, given his acceptance of the GRC, this is the identity function : a thought concerns the context of its tokening. The context, hence the topic situation which is the value of the function, varies, but the function is fixed by the system and does not vary. This is an instance of what I earlier described as 'parametric invariance' (Chapter 33).

Parametric invariance

When a determinant of truth-value is invariant in the strong, absolute sense — as Z-land is in Perry's tale — it is represented only implicitly. It is not an aspect of the *lekton* but belongs to the situation which the *lekton* concerns. However, contrary to what Millikan suggests, absolute invariance is not a necessary condition for implicitness. Following Perry, I argued that parametric invariance suffices. So the place and time of perception are elements which vary systematically and whose variation determines a corresponding variation in the veridicality conditions of the perception : the perception is true if and only if its content holds *in the situation of perception*. For Millikan, that is enough to show that the time and place in question belong to the explicit content of the perceptual representation (since they are expressed as values of variables that can accept alternative values). Still, as Perry rightly points out, there is an element of invariance there : a perception can *only*

---

[99] Lewis's individuals are world-bound. They have counterparts in other worlds but exist only in their own world.

represent what is going on at the time and place of perception. The function (here, the identity function) from the context of the perceptual episode to the relevant aspect of its truth-conditions is fixed by the perceptual mode and cannot vary, even if its values vary systematically. This is sufficient to distinguish the values of the function in question, which are provided automatically (as a function of the time and place of the perceptual episode), from the elements of content that are explicitly represented and can vary freely, without being constrained in this way.

Parametric invariance as I have described it involves the following four elements : the mode of representation determines a function from the context of tokening to some implicit aspect of the representation's satisfaction conditions. (We can perhaps treat absolute invariance as a special case, where the function is a constant function.) This framework makes clear that the GRC holds only in special cases, namely, the cases in which the function is the identity function. Perception is the typical example: the situation a perception concerns is the situation of perception, so the (Generalized) Reflexive Constraint holds (Chapter 36). But, as we have seen, the GRC does not hold in the case of memory. Whereas a perception represents what is going on at the time and place of perception, an episodic memory represents a past event which the subject has experienced. In both cases, the mode fixes the situation which the representation concerns as a function of the situation in which the perceptual or memory episode takes place ; but the function is identity only in the case of perception.

In memory, as in all the cases in which the GRC fails, the index of evaluation shifts away from the index of the context. Rather than view this as a special case to be accounted for by positing an implicit operator that is absent in the normal case, I have just suggested that we view the cases in which the GRC holds as the special cases in which the function determined by the mode is the identity function. On this account the shift, or the absence of shift, is *built into the function that is fixed by the mode* ; so the cases in which the index shifts and the cases in which it does not are exactly on a par. Whenever the function is not the identity function, the index of evaluation 'shifts'. This sort of shift ('*m*-shift', where '*m*' stands for 'mode') is very different from that which is effected through the use of operators ('*o*-shift', where 'o' stands for 'operator'). There is a clear sense in which an *o*-shift represents a departure from the normal/current situation of evaluation. But, arguably, an *m*-shift constitutes no such departure.

In instances of *o*-shift, the representation is syntactically complex. It involves an operator and an embedded representation in the scope of that operator. Like the complete representation, the embedded representation possesses a content. Both the content of the complete representation and that of the embedded representation can be modeled as functions from situations to truth-values. Following standard compositional procedure, the content of the complete representation will be said to map a situation s to the true iff for Q situation s' such that s' bears a certain relation R to s, the content of the embedded representation maps s' to the true. (The operator at stake will determine both the nature of the determiner Q and the nature of the relation R.) So, in the course of evaluating the complex representation with respect to a given situation s, one shifts from s to s' and one evaluates the embedded representation with respect to s'. Since s is the current situation of evaluation for the complete representation, an *o*-shift indeed represents a departure from the current situation of evaluation.

In addition to its content, the complete representation has a mode, which determines the situation s with respect to which it is to be evaluated. That situation, as we have seen, is a function of the situation or context c in which the representation is tokened. Whenever s ≠ c, that is, whenever the function determined by the mode is not the identity function, an *m*-shift takes place. Such a shift from c to s is very different from an *o*-shift, since an *o*-shift is a shift from s (the global situation of evaluation for the complete representation) to s' (the shifted situation of evaluation for the embedded representation). The two sorts of shift are so different that it may be a mistake to use the same word 'shift' in both cases. (The use of that single word was justified in the case of Lewis by his belief that there is a single phenomenon : for him, every failure of the GRC has to be accounted for in terms of an *o*-shift.)

Chapter 40
Free shiftability

Two features : automaticity and (generalized) reflexivity

In Chapters 32-33, I rejected the Externality Principle, on the grounds that the situation a representation concerns need not be determined by the external environment. A representation may concern a situation that the subject is thinking or talking about, rather than the situation he or she is in.[100] So for example the sentence 'At this time of the year it is freezing' may be uttered in reference to Chicago, the place being spoken about, even though the place of utterance is Berkeley. I described such cases as cases of 'free shiftability' (as opposed to 'operator-shiftability').

Free shiftability, as illustrated by this sort of example, has something in common with *m*-shiftability : it is a counterexample to the Generalized Reflexive Constraint. The speaker is in Berkeley (place of utterance) but the place relevant to the evaluation of what is said is Chicago. Similarly, the situation a memory concerns is temporally distinct from (anterior to) the situation the subject who remembers is in. In both cases, the situation the representation concerns is distinct from the situation in which the representation is tokened. The evaluation index is not the index of the context.

Still, there is also an important difference between cases like 'It is freezing' and memory. In memory, just as in perception, the psychological mode automatically determines the situation of evaluation as a function of the context, and it leaves the thinker no choice : a perception can only concern the situation the subject is currently experiencing, and a memory can only concern a certain situation she has

---

[100] « Human thought-processes differ radically from the analogous processes in animals, and, in particular, by their capacity to be detached from present activity and circumstances. Our thoughts may float free of the environment : we may follow a train of thought quite irrelevant to our surroundings or what we are engaged in doing. » (Dummett 1991 : 285)

experienced in the past, namely, that which left the memory as a trace. This is what I described as parametric invariance. But free shiftability is not so constrained (it seems). When I say 'It is freezing', I am free to talk of whatever place I like. It may be the place of the context, but it may be any other place that I happen to be talking about.

The freedom we have to shift the topic situation as discourse proceeds suggests that there is no 'invariance' at all in this sort of case, whether absolute or parametric. In the cases of parametric invariance I have talked about so far, the situation of evaluation is determined automatically as a function of the context ; it is highly constrained, even though it varies systematically as the context itself varies. But in discourse the situation of evaluation is not so constrained. Not only does it vary, it varies freely. The only thing that matters is the hearer's 'uptake', that is, his or her recognition of the speaker's intention to talk about this or that situation, however remote and different from the situation of utterance.

Table 3 summarizes the commonalities and differences we have just registered :

	perception	memory	discourse
Generalized reflexivity ?	*yes*	*no*	*it depends*
Automaticity ?	*yes*	*yes*	*no*

Table 3

As the table shows, perception and memory are two cases of parametric invariance for which the situation of evaluation is determined automatically as a function of the situation in which the representation is tokened. In discourse there is no such automaticity : what determines the topic situation is the speaker's intention (provided it is made sufficiently manifest to the hearer). At the same time, perception is distinguished from *both* memory and discourse by the fact that the situation of evaluation is bound to be the situation the subject is in when entertaining the representation : the Generalized Reflexive Constraint is satisfied, while it is not satisfied in the case of memory. In the case of discourse the GRC is not satisfied

either, since the topic situation may be very different from the situation of utterance. However, nothing prevents the topic situation from being the situation of utterance, in the discourse case, while this is simply not possible in the case of memory. This latitude, expressed by the response 'it depends' in the slot for generalized reflexivity, is undoubtedly linked to the freedom with which the situation of evaluation can be shifted (or not) in discourse, as opposed to the automaticity/rigidity of its determination in standard cases of parametric invariance.

Reducing free shiftability to o-*shiftability*

Operator-shiftability has been opposed both to free shiftability, the most salient examples of which can be found in discourse, and to *m*-shiftability, the most salient examples of which are provided by certain psychological modes such as memory. But we have just seen that there are also differences between free shiftability and *m*-shiftability. Does it mean that we end up with three distinct categories of shift ? Not necessarily. Despite the superficial differences between them, we can attempt to reduce free shiftability to one of the other two categories. There are two ways to do this, of course. We can (attempt to) reduce free shiftability to operator-shiftability or to *m*-shiftability. In this section, I will be concerned with the first of these options. The second option, which I favour, will be presented next.

If we use the invariance criterion which both Perry and Millikan appeal to, the apparently unconstrained freedom we have to shift the situation in discourse suggests that, even though it is tacitly rather than explicitly referred to, still the topic situation in such cases belongs with those elements of content that are explicitly represented as opposed to those that are carried implicitly, via the mode of representation.

It seems odd to say that something implicit can be part of explicit content, but the idea is simply that of an 'unarticulated constituent of the *lekton*', talked about at length (and eventually rejected) in Part IX. There we discussed the case of Perry's son saying 'It is raining' after a phone call to his brother, and meaning by this that it is raining *at the brother's place* (Murdock). Perry's son tacitly referred to his brother's place, and the question was : Is the place in question the place which the utterance 'concerns', or is it a constituent of the *lekton* ? Perry's own tendency is to treat this sort of unarticulated constituent as a constituent of the *lekton*, unarticulatedness

notwithstanding. Similarly, we might say that the place (viz. Chicago) is a constituent of the *lekton* in 'At this time of the year, it is freezing' because it lacks the right sort of invariance and can be shifted freely.

The suggestion, then, is that 'It is raining' is elliptical for 'It is raining in Murdock' in Perry's example, and that 'At this time of the year it is freezing' is likewise elliptical for 'In Chicago, at this time of the year, it is freezing' in the other example. 'Elliptical' here should not be understood in the strict syntactic sense, but in a more general, semantic sense.[101] What the subject means, the belief he expresses, is that it is raining/freezing *in Murdock/Chicago*. But some constituent of the content thereby expressed is left unarticulated because it is contextually inferable. Despite being unarticulated, the constituent in question is a constituent *of content*. To show this it is tempting to appeal to the 'mental articulation argument' which I discussed (and criticized) in Chapter 32 :

1. In this context the utterance 'It is raining' expresses the speaker's belief that it is raining in Murdock.
2. The place is an explicit constituent of (the content of) the speaker's belief that it is raining in Murdock. In other words : the subject is thinking about Murdock, i.e. *entertaining a mental representation of Murdock*.
3. An utterance has the same content as the belief it expresses.
4. Therefore, the place (Murdock) is a constituent of the content of the speaker's utterance just as it is a constituent of the content of his belief. The only difference is that it is linguistically unarticulated while it is mentally articulated.

This argument I tentatively reconstructed on the basis of Perry's description of the Murdock case. In criticizing it, I argued that the speaker's belief itself may 'concern' Murdock rather than being explicitly about it. Because that option has not been ruled

---

[101] See Neale 2000 : 286-87 on the 'sellarsian' sense of 'elliptical'. The general idea is that of 'free enrichment' : the sentence is given a richer interpretation in context than its literal meaning strictly licenses. In the case of implicit operators like 'in Murdock' or 'in Chicago', see Recanati 2002 for a description of the type of enrichment at issue.

out, the argument is without force. But the fact that the place *need not* be construed as a constituent of content does not entail that it *cannot* be so construed. On the view I am presently discussing, the place is an unarticulated constituent of the *lekton*. Even though the sentence is a simple sentence ('It is raining'), the content it contextually expresses — the *lekton* — has a more complex structure, due to the presence of the implicit operator 'In Murdock'.

When, on behalf of Lewis, I introduced the idea of an implicit operator to handle alleged counterexamples to the GRC, I rejected it on the grounds that this strategy has unwelcome consequences and leads one to put too much into the *lekton* (Chapter 38). But this criticism was targeted at the use of the implicit operator strategy to deal with instances of *m*-shift. Now we are talking about free shiftability in discourse, and the implicit operator strategy might well be legitimate in this particular case, even if it is objectionable as a way of dealing with *m*-shifts.

Reducing free shiftability to m*-shiftability*

In Chapter 33 I rejected the very idea of unarticulated constituents of the *lekton*, and I suggested that the notion of parametric invariance could be extended so as to cover the Murdock case. This suggestion is what I am now going to pursue.

Perry himself distinguishes several 'modes of thinking or discourse' about the weather, depending on whether one is talking about the local weather or about the weather at some other places. When one talks about the local weather, the relevant parameter is the place of utterance : one's weather talk concerns the place where one is (and the GRC is satisfied). Why not say that, in the other type of case, the weather statement concerns whichever place happens to be cognitively salient or is currently the topic of conversation ? On this view, the two sorts of cases are on a par, appearances notwithstanding — they are both cases of parametric invariance. In both cases, the mode of discourse fixes the relevant parameter (the place of utterance *vs* the place one is talking about). There is variation in both cases, and in both cases the variation is channeled through the relevant parameter.

What about the idea of free, unconstrained variability ? How can we explain the fact that the place of evaluation can be *any* place, without restriction, including

the place of utterance ? Doesn't this unrestrictedness show that there is no invariance whatsoever in that sort of case, not even parametric invariance ?

No, it does not. Whenever there is parametric invariance, there is something that does not vary (the function) and something that varies (the value of the function, which varies as a function of its arguments). If we decide to treat free shiftability as another case of parametric invariance, the relevant parameter will the topic of conversation. Now the value of that parameter varies as a function of the intentions which the speaker makes sufficiently manifest to the hearer. What varies freely here are the speaker's intentions, which are indeed fairly unconstrained. But those intentions serve as *argument* to the function, and the arguments to the function are unconstrained in *all* cases of parametric invariance. When the mode of discourse concerns local weather, the argument is the place of utterance, and that is unconstrained too : one can say 'It is raining' *anywhere*, and wherever one says it (in that mode), what one says will be true iff it is raining there. Similarly, in the other mode, one can intend to characterize the weather at any place by saying 'It is raining', and whatever the place one thereby intends to characterize, what one says will be true iff it is raining in that place.

Free shiftability is not to be found only in discourse. Our thoughts too may concern whatever entity is currently salient in our mind, even though the entity in question is not articulated in the thought that concerns it, but only in some other mental representation that serves as cognitive background for it. To accept this relation between thoughts, where one thought serves as background for another and determines what the other concerns, is to accept the existence of a mode characterized by the parameter : *thing currently being talked or thought about.* This mode, whether in discourse or thought, we may call the anaphoric mode. It is involved in the phenomenon I described as 'mental projection' in earlier writings (Recanati 1997, 2000b):

> A situation is first mentioned (that is, it is a constituent of some fact which is stated) ; the speaker then 'projects herself in that situation' or assumes it, and states something with respect to that situation, as if it was given in the external context (while it is only given 'in the discourse'). (Recanati 2000b : 67)

The idea that in the anaphoric mode one does 'as if' something was given in the context, while it is given only mentally or discursively, will be pursued in Chapter 41.

Chapter 41

The anaphoric mode :
a Bühlerian perspective

The primacy of the egocentric

Egocentrism is the view that every thought is about the *hic et nunc* situation — about the context of its tokening. On this view, a simple sentence like 'It is raining' talks about the place and time of the context, *unless* it is elliptical for a more complex sentence, 'At *l*, it is raining'. Such a sentence headed by an operator is *also* evaluated with respect to the current context, even though, in the course of so evaluating it, the embedded sentence it contains is evaluated with respect to a shifted index.

Against this view, I argued that only certain modes (e.g. the perceptual mode) are egocentric and satisfy the GRC. Other modes (e.g. the memory mode) impose a shift away from the index of the context. Such a shift, imposed by the mode, is what I called an *m*-shift. In Chapter 40, I suggested that free shiftability itself results from an *m*-shift. Through such a shift, a simple sentence like 'It is raining' *can* be used to talk about a different place than the place of the context, without being prefixed with an implicit operator. To achieve that result, we need not make the sentence more complex through the addition of implicit material ; we only have to *use* it in a different way, or in a different mode. Thus I distinguished the anaphoric mode from the egocentric mode. A representation in the anaphoric mode concerns any situation currently in focus, whether or not it is the situation in which the representation itself is tokened.

Egocentrism holds that there is only one mode or family of modes: the egocentric mode(s). That is what I reject. But that does not mean that I cannot accept a certain primacy of the egocentric. It is the universality of the egocentric mode (or family of modes) that I reject, but that is compatible with acknowledging its primacy.

The primacy I have in mind is entailed by views according to which one can go beyond egocentricity and entertain thoughts concerning situations distinct from the *hic et nunc* situation only by indulging in a form of mental simulation. Thus, according to Karl Bühler, we free ourselves from the *hic et nunc* situation by an act of 'imaginative displacement' : in the anaphoric mode, we treat a remote situation which has been evoked *as if* it was egocentrically given.[102] Talking about a situation indeed makes it 'present to the mind' in such a way that it can serve as a substitute for the egocentric situation and give rise to 'imagination-oriented deixis'. As Bühler writes, « remembered or imagined situations with a character similar to perception come up and replace the primary givenness of perceived situations » (Bühler 1934 : 149). On this view the anaphoric mode evolved from a prior capacity to simulatively shift the egocentric situation, hence it presupposes the egocentric mode, whose primacy must be acknowledged.

Cognitive significance again

In Chapter 32 I mentioned an argument against the view that, in the Murdock case, the utterance 'It is raining' *concerns* Murdock and expresses the proposition (*lekton*) that it is raining, rather than the proposition that it is raining in Murdock. The argument, due to Michael Devitt, runs as follows. When someone believes that it is raining and does not want to get wet, he or she typically does certain things. But in the Murdock case, the relevant dispositions to action are absent. After his phone call to his brother in Murdock, Perry's son says that it is raining, but he is not disposed to grab his umbrella. It seems, therefore, that he does *not* believe the same thing when he says 'It is raining' after the phone call to Perry's older son in Murdock, as he expresses when he says 'It is raining' after looking out the window in Palo Alto. In the latter case, arguably, he harbors the same belief as anyone who looks out the window and sees rain ; and the content of that belief can be represented as the place-relative proposition that it is raining. But in the Murdock case what he believes is the complete proposition that it is raining *in Murdock*. That explains why, in the

---

[102] That simulative process is what Campbell calls 'deep decentring' (Campbell 2002 : 183-87).

Murdock case, he is not disposed to act in the way someone who believes that it is raining does: by grabbing one's umbrella, cancelling out picnic plans, and so on and so forth.

Still, I want to maintain that even in the Murdock case, Murdock is not a constituent of the *lekton* ; it is only the place which the utterance concerns. How, then, can we account for the fact that the cognitive significance is different (as shown by the different dispositions to act) in the Murdock case and in the local weather case ? If the *lekton* is the same, how can there be a difference in cognitive significance ?

It will not do to invoke the other component of the Austinian proposition, namely the situation which the utterance concerns. For the brother in Murdock also believes that it is raining, and *he* is disposed to act in the relevant way. Perry's son, when talking to him on the phone from Palo Alto, has no such dispositions, even though he believes the same thing *concerning the same place* (Murdock). On this description, Perry's son in Palo Alto and his brother in Murdock have beliefs the complete contents of which can both be represented as the Austinian proposition consisting of the situation Murdock and the place-relative proposition that it is raining. Since the Austinian proposition is the same, while the cognitive significance is different, this seems to raise a problem for the SMR framework. (This version of the objection is due to Isidora Stojanovic.)

In reply, I would say that the difference is to be accounted for in terms of mode. What determines Murdock as the situation of evaluation for the place-relative proposition is the egocentric mode in one case (when the brother in Murdock says 'It is raining') and the anaphoric mode in the other case (when Perry's son in Palo Alto says 'It's raining' right after his conversation on the phone with his brother). The dispositions to act that are normally associated with the proposition that it is raining are *inhibited* when the *lekton* is entertained in the anaphoric mode. This comports well with the idea that the anaphoric mode rests on some form of simulation, since the simulation of mental states is generally seen as involving inhibition of the associated dispositions to act.

o-*shiftability and the anaphoric mode*

In Chapter 38, to illustrate the anaphoric mode, I used another meteorological example : 'At this time of the year, it is freezing'. This is supposed to be said about a place (Chicago) that has been evoked in previous discourse, as in the following sequence:

I would not like to live in Chicago. At this time of the year, it is freezing.

The first sentence in the sequence mentions Chicago. The second sentence is evaluated *with respect to* Chicago. This is like the example of projection I have already mentioned (Chapter 31) :

Berkeley is a nice place. There are bookstores and coffee shops everywhere.

Here too the second sentence is evaluated with respect to a location mentioned in the first sentence (and the quantifier rangers over bookstores and coffee shops at that location).

According to Bühler, we can also *incorporate* into the sentence the material upon which it is anaphoric — the 'exposition', as he calls it (Bühler 1934 : 426-7). Instead of a two-sentence discourse, we get :

In Chicago, at this time of the year it is freezing.

In Berkeley, there are bookstores and coffee shops everywhere.

The sentence in the scope of the operator 'in Chicago' or 'in Berkeley' is still in the anaphoric mode, but the material it is anaphoric to is now prefixed to the sentence instead of being external to it. The integrated sentences that result from incorporation bear the same relation to their two-sentence counterparts as conditionals do to what, in *Oratio Obliqua, Oratio Recta* (Recanati 2000b), I called 'Ducrot-Mackie pairs' :

Ducrot-Mackie pair :

Suppose he comes late. We won't worry him by mentioning the fire.

Conditional :

If he comes late, we won't worry him by mentioning the fire.

In the Ducrot-Mackie pair, the first sentence introduces a hypothetical situation, and the second sentence, in the anaphoric mode, is evaluated with respect to that situation. In the corresponding conditional, the antecedent is incorporated into the sentence, which now consists of two parts, but the relation between the two parts is much the same.[103]

<u>Egocentric mode</u>

simulation
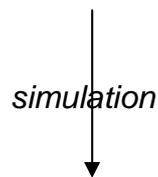
<u>Anaphoric mode</u>
(free shiftability)

incorporation
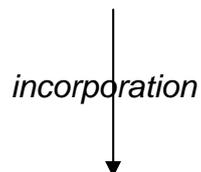
<u>*o*-shiftability</u>

Figure 2 : The Bülherian picture

---

[103] In Campbell's terms, that means that even in the integrated sentences headed by an operator, 'deep decentring' is at work. Campbell himself applies that idea to the analysis of tensed thoughts (Campbell 2002 : 186-87).

The benefits of incorporation have been insightfully described by Cosmides and Tooby within an evolutionary framework. As soon as one departs from the egocentric stance, that is, as soon as one is able to entertain representations that do not concern the *hic et nunc* situation, one needs a way of 'tagging' these representations so as to relativize their content to the right 'envelope of valid conditions' (Cosmides and Tooby 2000 : 105). What Cosmides and Tooby call 'scope-syntax', i.e. a recursive system of structured representations consisting of an operator and an embedded representation in the scope of the operator,[104] has evolved for this purpose, they claim. Be that as it may, Bühler's incorporation idea suggests that, far from reducing to *o*-shiftability, as both Lewis's and Perry's views imply, free shiftability is more basic. We have seen that it can be construed as a form of *m*-shiftability, involving the anaphoric mode. If Bülher is right, *o*-shiftability itself exploits that mode : it presupposes free shiftability, rather than being presupposed by it (Figure 3). On the Bühlerian picture, the anaphoric mode, which makes free shiftability possible, is accounted for in terms of the egocentric mode plus simulation, and *o*-shiftability is accounted for in terms of free shiftability plus incorporation.

---

[104] Such complex representations are what, in Recanati 2000b, I call 'δ-structures'.

References

Abusch, D. (1997). Sequence of Tense and Temporal *De Re*. *Linguistics and Philosophy* 20: 1-50.

Austin, J.L. (1971), *Philosophical Papers*, 2nd ed. Oxford: Clarendon Press.

Bach, K. (forthcoming). Searle against the World : How can Experiences Find their Objects ? To appear in S. Tsohatzidis (ed.) *John Searle's Philosophy of Language : Force, Meaning, and Thought*. Cambridge : Cambridge University Press.

Bar-Hillel, Y. (1954). Indexical expressions. In Bar-Hillel 1970 : 69-88.

Bar-Hillel, Y. (1963). Remarks on Carnap's *Logical Syntax of Language*. In Bar-Hillel 1970 : 116-142.

Bar-Hillel, Y. (1970). *Aspects of Language*. Jerusalem: The Magnes Press.

Barwise, J. (1989a). Situations, Facts, and True Propositions. In Barwise 1989b : 221-54.

Barwise, J. (1989b). *The Situation in Logic*. Stanford : CSLI Publications.

Barwise, J. and J. Etchemendy (1987). *The Liar : An Essay on Truth and Circularity*. New York: Oxford University Press.

Barwise, J. and J. Perry (1983). *Situations and Attitudes*. Cambridge, Mass: MIT Press/Bradford Books.

Van Benthem, J. (1977). Tense Logic and Standard Logic. *Logique et Analyse* 80 : 395-437.

van Benthem, J. (1985). *Modal Logic and Classical Logic.* Naples : Bibliopolis.

Blackburn, P. (1994). Tense, Temporal Reference and Tense Logic. *Journal of Semantics* 11 : 83-101.

Blackburn, P. (2000). Representation, Reasoning, and Relational Structures : a Hybrid Logic Manifesto. *Logic Journal of the IGPL* 8 : 339-65.

Blackburn P., de Rijke M., and Y. Venema (2001). *Modal Logic*. Cambridge : Cambridge University Press.

Brännmark, J. (forthcoming). Speaking for the Community : On Morality and We-Judgments. *Ms.*

Brogaard, B. (forthcoming). *Transient Truths : An Essay on the Metaphysics of Propositions. Ms.*

Bühler, K. (1934). *Theory of Language*. Eng. trans. by D.F. Goodwin, Amsterdam: John Benjamins, 1990.

Burge, T. (1991). Vision and Intentional Content. In E. Lepore and R. Van Gulick (eds.), *John Searle and His Critics*, Oxford : Blackwell, pp. 195-213.

Campbell, J. (1999). Immunity to Error through Misidentification and the Meaning of a Referring Term. *Philosophical Topics* 26 : 89-104.

Campbell, J. (2002). *Reference and Consciousness*. Oxford : Clarendon Press.

Carnap, R. (1937). *The Logical Syntax of Language*. London: Routledge and Kegan Paul.

Castañeda, H.-N. (1999). *The Phenomeno-Logic of the I : Essays on Self-Consciousness* (edited by J. Hart and T. Kapitan). Bloomington : Indiana University Press.

Chauvier, S. (2001). *Dire « je » : Essai sur la subjectivité*. Paris : Vrin.

Chierchia, G. (1989). Anaphora and Attitudes *De Se*. In R. Bartsch, J. van Benthem & P. van Emde Boas (eds.) *Semantics and Contextual Expression*. Dordrecht : Foris, p. 1-31.

Corazza, E. (2004). *Reflecting the Mind : Indexicality and Quasi-Indexicality*. Oxford : Clarendon Press.

Cosmides, L. and Tooby, J. (2000). Consider the Source : the Evolution of Adaptations for Decoupling and Metarepresentations. In D. Sperber (ed.), *Metarepresentations : A Multidisciplinary Perspective*, Oxford : Oxford University Press, pp. 53-115.

Cresswell, M. (1990). *Entities and Indices*. Dordrecht : Kluwer.

Cresswell, M. (1996). *Semantic indexicality*. Dordrecht : Kluwer.

Dennett, D. (1978). Where Am I ? In Dennett, D. *Brainstorms*, Cambridge, Mass. : MIT Press/Bradford Books, pp. 310-23.

Devlin, K. (1991). *Logic and Information*. Cambridge : Cambridge University Press.

Donnellan, K. (1966). Reference and Definite Descriptions. *Philosophical Review* 75 : 281-304.

Ducrot, O. (1972). *Dire et ne pas Dire*. Paris: Hermann.

Dummett, M. (1991). Thought and Perception : The Views of Two Philosophical Innovators. In Dummett, M. *Frege and Other Philosophers*, Oxford : Clarendon Press, pp. 263-88.

Dummett, M. (1993). Testimony and Memory. In Dummett, M. *The Seas of Language*, Oxford: Oxford University Press, pp. 411-28.

Dummett, M. (forthcoming). Existence, Possibility and Time. *Ms.*

Egan, A. (forthcoming). Epistemic Modals, Relativism and Assertion. *Ms.*

Evans, Gareth (1982). *The Varieties of Reference*, ed. J. McDowell. Oxford: Clarendon Press.

Evans, G. (1985). Does Tense Logic Rest on a Mistake? In Evans, G. *Collected Papers*, Oxford: Clarendon Press, pp. 343-363.

Feldman, R. (1980). Saying Different Things. *Philosophical Studies* 38 : 79-84.

Fernandez, J. (forthcoming). The Intentionality of Memory. To appear in *Australasian Journal of Philosophy*.

Fintel, K.v. and Gillies, A. (2004). « Might » Made Right. *Ms.*

Fintel, K.v. and Gillies, A. (2006). CIA leaks. *Ms.*

Frege, G. (1918-19). The Thought: a Logical Enquiry. English trans. by A. and M. Quinton, in P. Strawson (ed.), *Philosophical Logic*, Oxford University Press (1967), 17-38.

Frege, G. (1967). *Kleine Schriften*, ed. I. Angelelli. Hildesheim : Olms.

Frege, G. (1970). *Translations from the Philosophical Writings of Gottlob Frege,* ed. P. Geach and M. Black. Oxford: Basil Blackwell.

Garcia-Carpintero, M. (1997). Indexicals as Token-Reflexives. *Mind* 106: 529-63.

Garcia-Carpintero, M. (forthcoming). Relativism, Vagueness, and What is Said. In M. Kölbel and M. Garcia-Carpintero (eds.), *Relativizing Utterance Truth*.

Geach, P. (1949). Review of J. Weinberg, *Nicolaus of Autricourt : A Study in Fourteenth Century Thought. Mind* 58 : 238-45.

Geach, P. (1955). Review of Mates 1953. *Philosophical Review* 64 : 143-5.

Geach, P. (1957). On Beliefs about Oneself. *Analysis* 18 : 23-24.

Goodman N. (1951). *The Structure of Appearance.* Third edition, Dordrecht : Reidel, 1977.

Harman, G. (1986). *Change in View.* Cambridge, Mass. : MIT/Bradford Books.

Higginbotham, J. (1995). Tensed Thoughts. *Mind and Language* 10 : 226-249.

Higginbotham, J. (2003a). Remembering, Imagining, and the First Person. In A. Barber (ed.), *Epistemology of Language*, Oxford : Oxford University Press, pp. 496-533.

Higginbotham, J. (2003b). Tensed Second Thoughts. In Jokic, A. and Smith, Q. (eds.), *Time, Tense and Reference*, Cambridge, Mass : MIT Press/Bradford Books, pp.

Hintikka, J. (1973). *Time and Necessity : Studies in Aristotle's Theory of Modality*. Oxford : Clarendon Press.

Hoerl, C. (2001). The Phenomenology of Episodic Recall. In C. Hoerl and T. McCormack (eds.), *Time and Memory*, Oxford : Oxford University Press, pp. 315-335.

James, W. (1890). *Principles of Psychology*, vol. 1. London : Macmillan.

Kamp, H. (1971). Formal Properties of 'Now'. *Theoria* 37 : 227-73.

Kaplan, D. (1989). Demonstratives. In Almog, J., Perry, J. and Wettstein, H. (eds.), *Themes from Kaplan*, New York : Oxford University Press, pp. 481-563.

Katz, J. (1977). *Propositional Structure and Illocutionary Force*. New York: Crowell.

King, J. (2003). Tense, Modality, and Semantic Values. *Philosophical Perspectives* 17 : 195-245.

Kölbel, M. (2002). *Truth Without Objectivity*. London : Routledge.

Kölbel, M. (forthcoming). Increasing the Adicity of a Truth-Predicate. In M. Kölbel and M. Garcia-Carpintero (eds.), *Relativizing Utterance Truth*.

Kripke, S. (1980), *Naming and Necessity.* Oxford: Blackwell.

Kuhn, S. (1980). Quantifiers as Modal Operators. *Studia Logica* 39 : 145-58.

Landman, F. (1991). *Structures for Semantics*. Dordrecht : Kluwer.

Lasersohn, P. (1999). Pragmatic Halos. *Language* 75 : 522-51.

Lasersohn, P. (2005). Context Dependence, Disagreement, and Predicates of Personal Taste. *Linguistics and Philosophy* 28: 643–686*.*

Lewis, D. (1970). General Semantics. *Synthese* 22 : 18-67, reprinted (with a postscript) in Lewis 1983, pp. 189-232.

Lewis, D. (1978). Truth in Fiction. *American Philosophical Quarterly* 15 : 37-46, reprinted (with a postscript) in Lewis 1983, pp. 261-280.

Lewis, D. (1979). Attitudes *de dicto* and *de se. Philosophical Review* 88 : 513-43, reprinted (with a postscript) in Lewis 1983, pp. 133-159.

Lewis, D. (1980). Index, context, and content. In S. Kanger and S. Öhman (eds.), *Philosophy and Grammar,* Dordrecht: Reidel, pp. 79-100. Reprinted in Lewis 1998, pp. 21-44.

Lewis, D. (1983). *Philosophical Papers, Volume 1*. New York: Oxford University Press.

Lewis, D. (1998). *Papers in Philosophical Logic*. Cambridge : Cambridge University Press.

Loar, B. (1976). The Semantics of Singular Terms. *Philosophical Studies* 30 : 353-377.

Ludwig, P. (2005). Variétés de la connaissance de soi. In S. Bourgeois-Gironde (ed.) *Les Formes de l'indexicalité*, Paris : Editions Rue d'Ulm, pp. 95-117.

MacBride, F. (forthcoming). Predicate Reference. In E. Lepore & B. Smith (eds.) *The Oxford Handbook of Philosophy of Language* (Oxford University Press, 2006).

MacFarlane, J. (2003). Future Contingents and Relative Truth. The *Philosophical Quarterly* 53 : 321-336.

MacFarlane, J. (forthcoming). Nonindexical Contextualism. *Ms.*

Mackie, J. (1973). *Truth, Probability, and Paradox*. Oxford: Clarendon Press.

Malmberg, B. (1979). *Le Langage, signe de l'humain*. Paris : Picard

Martin, M. (2001). Out of the Past : Episodic Recall as Retained Acquaintance. In C. Hoerl and T. McCormack (eds.), *Time and Memory*, Oxford : Oxford University Press, pp. 257-284.

Mates, B. (1953). *Stoic Logic*. University of California Press.

McGinn, C. (1982). The Structure of Content. In Woodfield, A. (ed.), *Thought and Object*, Oxford : Clarendon Press, 207-58.

Mellor, H. (1998). *Real Time II*. Cambridge : Cambridge University Press.

Millikan, R. (2004). *Varieties of Meaning*. Cambridge, Mass. : MIT Press/Bradford Books.

Millikan, R. (2006). Indexicals, Demonstratives and the Explanation of Behavior. *Ms.*

Nagel, T. (1970). *The Possibility of Altruism*. Princeton, N.J. : Princeton University Press.

Neale, S. (2000) On Being Explicit : Comments on Stanley and Szabo, and on Bach. *Mind and Language* 15 : 284-94.

Nozick, R. (1981). *Philosophical Explanations*. Cambridge, Mass. : Harvard University Press.

Ogihara, T. (1996). *Tense, Attitudes, and Scope*. Dordrecht : Kluwer.

Parsons, T. (1990). *Events in the Semantics of English*. Cambridge, Mass. : MIT Press.

Partee, B. (1973). Some Structural Analogies between Tenses and Pronouns in English. *Journal of Philosophy* 70 : 601-7, reprinted in Partee 2004, 50-58.

Partee, B. (2004) *Compositionality in Formal Semantics*. Oxford : Blackwell.

Peacocke, C. (1981). Demonstrative Thoughts and Psychological Explanation. *Synthese* 49: 187-217.

Peacocke, C. (1983). *Sense and Content*. Oxford: Clarendon Press.

Perry, J. (1977). Frege on Demonstratives. *Philosophical Review* 86 : 474-97, reprinted, with a postscript, in Perry 1993b pp. 3-32.

Perry, J. (1979). The Problem of the Essential Indexical. *Noûs* 13 : 3-21, reprinted, with a postscript, in Perry 1993b pp. 33-52.

Perry, J. (1980). Belief and Acceptance. *Midwest Studies in Philosophy* 5 : 533-42, reprinted, with a postscript, in Perry 1993b pp. 53-67.

Perry, J. (1986a). Perception, Action, and the Structure of Believing. In R. Grandy and R. Warner (eds.) *Philosophical Grounds of Rationality*, Oxford : Clarendon Press, pp. 333-61 ; reprinted in Perry 1993b, pp. 121-49.

Perry, J. (1986b). Thought without Representation. *Proceedings of the Aristotelian Society, Supplementary Volume* 60: 137-51, reprinted in Perry 1993b, pp. 205-19.

Perry, J. (1993a). Postscript to 'Thought without Representation'. In Perry 1993b, pp. 219-25.

Perry, J. (1993b). *The Problem of the Essential Indexical and Other Essays*. New York : Oxford University Press.

Perry, J. (2000). Myself and 'I'. In Perry, J. *The Problem of the Essential Indexical and Other Essays*, expanded edition, Stanford : CSLI Publications, pp. 325-339.

Perry, J. (2001). *Knowledge, Possibility and Consciousness*. Cambridge, Mass : MIT Press/Bradford Books.

Prior, A. (1957). *Time and Modality*. Oxford : Clarendon Press.

Prior, A. (1959). Thank Goodness That's Over. *Philosophy* 34 : 12-17, reprinted in Prior 1976 : 78-84.

Prior, A. (1967). *Past, Present and Future*. Oxford : Clarendon Press.

Prior, A. (1968a). 'Now'. *Noûs* 1 : 101-19, reprinted in Prior 2003 : 171-93.

Prior, A. (1968b). *Papers on Time and Tense*. Oxford : Clarendon Press.

Prior, A. (1971). *Objects of Thought*, ed. by P. Geach and A. Kenny. Oxford : Clarendon Press.

Prior, A. (1976). *Papers in Logic and Ethics*. London: Duckworth.

Prior, A. (1977), with K. Fine. *Worlds, Times and Selves*. London : Duckworth.

Prior, A. (2003). *Papers on Time and Tense*, second edition, ed. by P. Hasle, P.
Ohrstrom, T. Braüner and J. Copeland. Oxford : Clarendon Press.

Quine, W. v. O. (1960). Variables Explained Away. In his *Selected Logic Papers*,
Cambridge : Harvard University Press, 1995, 227-35.

Recanati, F. (1981). *Les Enoncés performatifs*. Paris : Editions de Minuit.

Recanati, F. (1987). *Meaning and Force*. Cambridge : Cambridge University Press.

Recanati, F. (1993). *Direct Reference: From Language to Thought*. Oxford:
Blackwell.

Recanati, F. (1995a). Le Présent épistolaire : une perspective cognitive.
L'Information Grammaticale 66 : 38-44.

Recanati, F. (1995b). Quasi-Singular Propositions: the Semantics of Belief Reports.
Proceedings of the Aristotelian Society, Supplementary volume 69 :175-194.

Recanati, F. (1996). Domains of Discourse. *Linguistics and Philosophy* 19 : 445-75.

Recanati, F. (1997). The Dynamics of Situations. *European Review of Philosophy* 2:
41-75

Recanati, F. (1999). Situations and the Structure of Content. In K. Murasugi and R.
Stainton (eds.), *Philosophy and Linguistics*, Westview, Boulder, Colorado, pp. 113-
165.

Recanati, F. (2000a). Opacity and the Attitudes. In P. Kotatko et A. Orenstein (eds.),
Knowledge, Language and Logic: Questions for Quine. Dordrecht : Kluwer, 2000, p.
363-402

Recanati, F. (2000b). *Oratio Obliqua, Oratio Recta. An Essay on Metarepresentation.*
Cambridge, Mass : MIT Press/Bradford Book.

Recanati, F. (2002). Unarticulated Constituents. *Linguistics and Philosophy* 25 : 299-
345.

Recanati, F. (2003). The Limits of Expressibility. In B. Smith (ed.), *John Searle*.
Cambridge: Cambridge University Press, pp. 189-213.

Recanati, F. (2004). *Literal Meaning*. Cambridge : Cambridge University Press.

Recanati, F. (2005a). Literalism and Contextualism : Some Varieties. In G. Preyer
and G. Peter (eds.) *Contextualism in Philosophy : Knowledge, Meaning, Truth*.
Oxford : Oxford University Press, pp. 171-196.

Recanati, F. (2005b). Loana dans le métro. Remarques sur l'indexicalité mentale. In S. Bourgeois-Gironde (ed.) *Les Formes de l'indexicalité*, Paris : Editions Rue d'Ulm, pp. 19-34.

Recanati, F. (2006a). Indexical Concepts and Compositionality. In M. Garcia-Carpintero et J. Macia (eds.) *Two-Dimensional Semantics*, Oxford : Oxford University Press, pp. 249-57.

Recanati, F. (2006b). Relativized Propositions. In M. O'Rourke and C. Washington (eds.), *Situating Semantics : Essays on the Work of John Perry*. Cambridge, Mass. : MIT Press/Bradford Books, pp. 000-000.

Recanati, F. (forthcoming a). It is raining (somewhere). To appear in *Linguistics and Philosophy*.

Recanati, F. (forthcoming b). Indexicality and Context-Shift. In Recanati, F. *Truth-Conditional Pragmatics*, Oxford : Clarendon Press.

Recanati, F. (forthcoming c). Millikan's theory of signs. To appear in *Phenomenology and Philosophical Research*.

Reichenbach, H. (1947). *Elements of Symbolic Logic*. London : Macmillan.

Richard, M. (1981). Temporalism and Eternalism. *Philosophical Studies* 39 : 1-13.

Richard, M. (1982). Tense, Propositions, and Meanings. *Philosophical Studies* 41 : 337-51.

Richard, M. (2003a). Introduction to Part I : The Philosophy of Tensed Language. In Jokic, A. and Smith, Q. (eds.), *Time, Tense and Reference*, Cambridge, Mass : MIT Press/Bradford Books, 25-45.

Richard, M. (2003b). Objects of Relief. In Jokic, A. and Smith, Q. (eds.), *Time, Tense and Reference*, Cambridge, Mass : MIT Press/Bradford Books,157-89.

Richard, M. (2004). Contextualism and Relativism. *Philosophical Studies* 119 : 215-42.

Salmon, N. (1986). *Frege's Puzzle*. Cambridge, Mass. : MIT Press/Bradford Books.

Sauerland, U. (2002). The Present Tense is Vacuous. *Snippets* 6, at `http://www.ledonline.it/snippets/`

Searle, J. (1969). *Speech Acts*. Cambridge: Cambridge University Press.

Searle, J. (1983). *Intentionality*. Cambridge: Cambridge University Press.

Searle, J. (1991). Reference and Intentionality. In E. Lepore and R. Van Gulick (eds.), *John Searle and His Critics*, Oxford : Blackwell, pp. 227-241.

Shoemaker, S. (1968). Self-Reference and Self-Awareness. *Journal of Philosophy* 65 : 555-67, reprinted in Shoemaker (2003), pp. 6-18.

Shoemaker, S. (1970). Persons and their Past. *American Philosophical Quarterly* 7 : 269-85, reprinted in Shoemaker (2003), pp. 19-48.

Shoemaker, S. (2003). *Identity, Cause, and Mind.* Expanded edition, Oxford : Clarendon Press.

Sperber, D. (1982). Comments on Clark and Carlson's Paper. In N. Smith (ed.), *Mutual Knowledge*, London : Academic Press, pp. 46-51.

Stalnaker, R. (1970). Pragmatics. *Synthese* 22 : 272-89.

Stalnaker, R. (1978). Assertion. *Syntax and Semantics* 9 : 315-32.

Stanley, J. and Z. Szabo (2000). On Quantifier Domain Restriction. *Mind and Language* 15 : 219-61.

Stojanovic, I. (forthcoming). Talking about Taste : Disagreement, Implicit Arguments and Relativized Contents. *Ms.*

Strawson, P. (1966). *The Bounds of Sense.* London : Methuen.

Tulving, E. (1993). What is Episodic Memory ? *Current Perspectives in Psychological Science* 2 : 67-70.

Vendler, Z. (1979) Vicarious Experience. *Revue de Métaphysique et de Morale* 84 : 161-73.

Vlach, F. (1973). *'Now' and 'Then' : A Formal Study in the Logic of Tense and Anaphora.* Ph. D, UCLA.

Walton, K. (1990). *Mimesis and Make-Believe.* Cambridge, Mass. : Harvard University Press.

Williams, B. (1973). Imagination and the Self. In his *Problems of the Self*, Cambridge : Cambridge University Press, pp. 26-45.

Wittgenstein, L. (1958). *The Blue and Brown Books.* Oxford : Blackwell.