

Context and Logical Form

Wolfram Hinzen

August 6, 2003

1 Introduction

Contextuality is a pervasive phenomenon in the study of cognition. Nonetheless, the way in which laws for the justification of induction, say, are context-dependent, might not relate in any direct manner to the way in which, say, meaning is contextual. Thus we won't be able to presuppose a general notion of context, or of the "problem of context". We will have to ask for each aspect of cognition separately what the notion and the role of context is.

This chapter analyzes contextuality as relevant to a science of meaning. I will first, in section 2, introduce a notion of meaning deprived of contextuality altogether. I then discuss some aspects of "dynamic" semantic theories, for which the contextuality of meaning is decisive (section 3), and argue that these models naturally lead to normative models of discourse, considered as a kind of human action. In section 4 I argue against Stalnaker's recent criticism of dynamic semantics and the representation of context he himself recommends. Finally, in section 5, I explore in what sense there might be a *linguistic* (sentence-internal) notion of context that is important to linguistic explanation.

2 Meaning as Non-Contextual

2.1 Meaning

How might meaning be contextual? I will take this to be a question about *linguistic* meaning, a notion that describes aspects of the interpretation of natural language expressions as determined by the rules of language themselves. Note it is thus, in particular, not a notion applying to features of *utterances*. By an *expression* I will mean a pairing of a sound and a meaning, and by a *language* a systematic way of generating such pairings.

The existence of a notion of linguistic meaning thus understood might be denied, for example on the ground that the rules of a language never determine any aspects of interpretation, or that they do it only together with non-linguistic factors. On such views, an interpretation accrues to a natural language sentence at least in part *externally*, for example by virtue of the fact that on an occasion of its use, a speaker of that language expresses certain beliefs or opinions with it. But there would be no interpretive aspects of expressions that are linguistic narrowly speaking and that we could isolate as such. But if linguistic aspects of meaning cannot be isolated from various non-linguistic ones that also determine aspects of meaning on an occasion of use, meaning is a rather messy phenomenon. How a speaker on an occasion of the use of an expression ascribes a certain interpretation to it depends on a myriad of factors and cognitive faculties. Scientific progress depends on suitable idealizations and theoretical terms that make lawful generalizations statable. These will usually depend on decontextualizing a phenomenon as complex as “interpreting an utterance in context”.

If, on the other hand, the claim is not only that linguistic aspects of interpretation cannot be isolated from non-linguistic ones, but that no aspects of interpretation *are* intrinsically linguistic, it is clearly false. For that in (1)

(1) [John expected to hit himself],

himself is necessarily (by the rules of language) the same as *John*, is clearly an interpretive property of (1). It follows from general aspects of its grammatical structure, in which *John* binds the reflexive. It is equally intrinsic to

(2) I wonder who [John expected to hit himself],

that in it *himself* necessarily isn't the same as *John*. That again follows from structural properties of it, namely that the reflexive is bound by a variable, which in turn is bound by *who*: a paraphrase of (2) might be

(2P) I wonder for which x , x a person, John expected x to hit x .

Note that these examples do not depend on a solution to the notoriously difficult question to free a notion of linguistic meaning from other ingredients of interpretation such as background beliefs, world knowledge, and other non-linguistic factors. I have no definition of “language” on offer here, except that I stipulate that underlying language use there is a system of grammatical knowledge represented in the mind. It seems clear enough from examples and actual work what a grammatical process or a rule of language is, and to the extent that the examples show such processes or rules to have interpretive effects, they give content to the notion of linguistic meaning.

It seems hard to question then, not only that there are interpretive aspects of expressions which are linguistic, but also that they are *non-contextual*. For the relevant interpretive aspects of (1) and (2) have nothing to do with the non-linguistic context in which the expressions are used. They are things that speakers of English *know*, being consequences of an abstract system of principles and rules. To the extent that the explanation of standard facts about analyticity, synonymy, or ambiguity that are known to a native speaker proceeds on the basis of such principles and rules, it would seem that non-linguistic context drops out from the theory of meaning. Whether a sentence is ambiguous, for example, would, if so explained, have nothing at all to do with how or why it is used, and how contextual the interpretations assigned to such uses may be. It would be explained by appeal to the fixed structure of an internally represented system of human knowledge. This then is the question we ask in this chapter: to what extent there *are* interpretive aspects of expressions which are linguistic but which *do* depend on the non-linguistic (epistemic, environmental, etc.) contexts in which they get used. I will discuss some recent views on which content and context become a unified phenomenon, and argue that on this view, meaning, far from being a linguistic notion, merges into a general notion of “speech act content”. The latter seems far too contextual to be theoretically tractable, on the other hand, shot through as it is by pragmatic inferential mechanisms for fixing it. Arising in part through general principles on how to act in a discourse context, it is open to normative description, much like human action in general. As such, it is not a notion on which we can build linguistics as a (naturalistic) science.

2.2 Logical Forms (LFs)

I will assume, in the tradition of generative grammar, that in a speaker’s mind/brain the interpretive properties of expressions are represented in a component of the language system that is standardly called *Logical Form* (LF). It is by appeal to the structure of the representation of an expression at that level of representation LF that we explain certain linguistic aspects of its meaning. To take a standard example, the sentence

(3) Pictures of everybody are on sale

can mean that everybody is such that pictures of him/her are on sale, or that pictures are on sale that have everybody on them (group pictures). But importantly, (3) cannot mean: for some pictures, x , everybody is such that x is a picture of him/her and x is on sale. How do we explain these empirical facts? In logic, we make the difference merely manifest, by associating two different logical forms with the expression, with appropriately different relative scope relations among its quantifiers. This is a *representation* of the relevant fact, not an explanation (and it doesn’t tell why the third reading is not possible). If we add that these different logical forms are associated ‘by convention’ to

the expression (3), we make no further step towards explanation: there are no theories of convention that tell us how this works, and it seems implausible in any case that highly constrained syntactic principles responsible in human languages for what scopes over what are the result of arbitrary conventions.

Generative grammar takes a different course, and attempts to *explain* the different logical forms associated with an expression by deriving them by syntactic means - mechanical processes like merging two expressions, targeting one syntactic structure so as take one constituent from it and insert it in a different place, and so on. It is by means of such processes that the quantified noun phrases *pictures* and *everybody* in (3) get their relative scope possibilities. The impossibility of the third reading is explained by the fact that its derivation violates independently motivated constraints on possible linguistic structures.¹

Note then that the notion of LF in its linguistic sense has little to do with logic, normative rules of inference, or commitments we undertake: each LF is a natural object in a speaker's mind/brain, and that it plays a role in logic and reasoning, if it plays such a role, is not what determines its properties.² It is important to note that the claim that there is such a thing as a component of LF in the human language system is an empirical claim that is right or wrong. By contrast, there is a sense in which a logical form in the sense of logic or the theory of reasoning comes into existence through its definition. It is clearly not a notion of natural language meaning, being a characterization of *thought*, which presumably is what Frege and Russell designed it for.³

For these reasons I will adopt a more recent label and speak not of logical forms, but of a component of *Semantic Representation* in the human faculty of language, SEM for short. To say then that there are interpretive aspects of expressions which are contextual would be to say that there are features of the extra-linguistic context that make linguistic expressions have certain properties at SEM. Considering its conception in recent grammatical theory ([Cho95]), there is a sense in which SEMs are not open to doubt. It would be the claim that grammatical operations make no contribution to semantic interpretation, language being pure sound. As argued above, this option seems mistaken. What might well be wrong is that SEM is the only interface at which the language system meets other cognitive components governing performance. That

¹See [Hua95], 132–135 for details.

²At best, it provides evidence for the linguist to ascribe certain properties to it.

³Normatively construed commitments as in Arló Costa's contribution (this volume) can have no role to play in a strongly naturalistic enterprise like the present one. The distinction between old Chomskyan 'competence' and Arló Costa's 'commitment', is crucial here, although both are contrasted with 'performance'. It will transpire further below, especially in the last part of this paper, that Arló Costa's proposed marriage between theories of linguistic meaning and normative models in formal epistemology should be made consistent with the fact that all aspects of meaning depending on the generative structure of expressions have nothing to do with the epistemological states of speakers.

however is an empirical question, however hard to answer, and for the moment the very best assumption one can make is the minimal one, according to which there is only one, and try to prove this assumption wrong. If so, SEM will not be negotiable at least for the moment.

I will suppose further that human languages consist of a list of primitive or unstructured *lexical items* (LIs) as distinct from a computational system (CS). A (mental) lexicon, on which the rules of the latter operate is a storage place for the sound and meaning properties of lexical items. Let these properties be expressed in terms of a collection of features. Each LI is a collection of *features* of different sorts, a representation of what a speaker knows about them, and it is these features which are fed into the computational component of language, which builds larger representations out of them.

These assumptions make the above contextual option rather unpalatable. Derivations in the faculty of language are encapsulated and have locality properties, first in the trivial sense that what the computations have access to are symbolic structures represented in the mind alone, and secondly in the sense of the empirical assumption that they proceed in the mind's linguistic component alone. On its way to compute a SEM, the CS doesn't have access to what is happening outside the head, or elsewhere in the head. What enters into a derivation is mental representations of a specific format that CS can compute, but non-linguistic context doesn't enter the computation of SEMs at all. This is more or less also the conclusion that I will reach below, even though there remain some interesting issues concerning the role of context variables *within* SEMs.

Jason Stanley has recently stated that "all effects of extra-linguistic context on the truth-conditions of an assertion are traceable to logical form" ([Sta00], 396). The sense of this claim will depend on clarifying, of course, the notion of logical form used. Stanley explains that it is the "real' structure" of a sentence, which may be "quite distinct from its surface grammatical form" (392). That leaves things rather open. In its generative-grammatical sense, in particular, the notion of LF is hardly describable in Stanley's terms, for whom it seems to depend on notions of Deep and Surface Structure long abandoned. The Minimalist Program (MP, [Cho95]) at least is also not coherent with Stanley's assumption that LF is a structural representation that is the object of compositional and referential semantic interpretation (cf. [Sta00], 399). While the notion of truth conditions is still informally used in MP, it is not *essential* to the theory of LFs that they determine such truth conditions. LFs are internally motivated, by the operations of a grammar. Truth conditions provide *data* for a theory of LF, but they are things to be *derived* on syntactic grounds (cf. the above example). Moreover, if informally used, the notion of truth conditions is, contrary to what Stanley assumes (e.g., 395), applied to *expressions*, which are objects, and not to assertions or utterances, which are acts.

Logical form being a highly theoretical term fixed by its definition in technical discourse, Stanley's main question will at least in part have an answer that simply follows from this definition. Thus if Stanley suggests that "all effects of extra-linguistic context

should be traced to logical form”, we might take him to suggest as theoretically useful a specific theoretical notion, rather than to answer a factual question. Contemplating its usefulness, it seems Stanley’s notion is very rich in terms of the information it captures. Nonetheless, there is some affinity in what he says to what I will myself claim in section 5. This is the idea that “the effects of context on the truth-conditional interpretation of an assertion are restricted to assigning the values to elements in the expression uttered” (396). If only this idea didn’t depend on the logical notions of “truth condition”, “semantic value” and “denotation” that seem questionable as notions of linguistic theory (on which more below), it seems congenial to the proposals to follow. That is to say: the role of context is the instantiation of certain variables occurring in SEM, and it doesn’t contribute to linguistic meaning in any other way. I will now delve into some of the ways in which the effects of extra-linguistic context on the interpretation of human speech is vast.

2.3 SEMs and Context

Context enters when SEMs are put to use by performance systems accessing the language system. The moment where language gets produced and interpreted, contextual effects abound. In a context the system of linguistic knowledge interfaces with human intentions, background assumptions, culture and custom, a perceptual situation, the climate, human temper, and more. Maybe someone using the word *sun* in the Sahara will use it with slightly different connotations than when he uses it in England. Maybe a great composer will say something slightly different from what we say when he says

(4) I like music.

Maybe he will *like* it in a different way than we do, and think of something quite different when he uses the word *music*, even though that might depend on whether he does it on a good or a bad day. All of this is possible, but if all of it is meant to be part of the *meaning* of what he says, that notion would appear to be rather inscrutable. We might put together the “meaning” of the word *music*, in this massively contextual sense, out of what he means and we mean and your friends mean and maybe what other populations mean, at certain times or others. Whatever we fabricate here would not seem to result in anything of theoretical interest.

Neither do the examples just put forward show anything about the contextuality of linguistic meaning. All of the contextuality just observed is perfectly consistent with the view that concepts are innate (and thereby non-contextual).⁴ On one such view, humans are born with a stock of concepts that get phoneticized upon exposure to an

⁴E.g. [Fod98] for one such view.

environment. Human nature apparently makes them similar enough to make communication reasonably effective, but that doesn't show anything about how much variance we may find across speakers and contexts, which may have different sources. The above observations could be redescribed thus: Even though, by way of innate endowment, the concept *music* is whatever it is, it just happens that contexts of language use embed the knowledge of *music* as specified in the mental lexicon of a speaker *A* in a rich background of opinions and feelings that *A* has about music in those contexts. Not *A*'s concepts vary, but his feelings. The same happens to speaker *B*. Some of the opinions and feelings of *A* and *B* may be shared, but that doesn't mean that beyond what's in the head of *A* and *B* there is some mysterious shared concept of music outside the heads of both (I return to this below).

Maybe in fact we have slightly different concepts in different contexts, which we associate to the same sound. But maybe not. The feature of human languages that Jerry Fodor calls their *systematicity* seems to depend on the fact that a word in one context means what it does in another. It's because music means the same in *I like music* and *I don't like music* that both sentences cannot be true together. Relative to an idiolect a word will have the interpretation it has, and it is with this interpretation that it enters syntactic compounds, to which it contributes precisely its interpretation, and nothing else. In each context of use it may happen that music is identified or valued in different ways. But what matters for lexical theory is the content a word contributes to the content of the compound in which it enters, and it doesn't seem to be the case that epistemic features of a context, changing criteria of identification, and so on, enter systematically in the formation of syntactic compounds, in which they seem invisible.

Thus it is everything but clear whether even upon adding a lexicon to a syntactic theory we open up our investigation to an intrusion of context that might be fatal to a naturalistic scientific endeavour regarding human language. That conclusion clearly depends on a choice of subject matter for linguistic theory, namely the idiolects of speakers (or I-languages, in Chomsky's term). For syntactic theory, that choice is perfectly standard, and I can't see how it could be negotiable for the theory of the lexicon. What would the units of linguistic computation be, on a more communication-theoretic view, say? Where the theory of the lexicon is part of a theory of I-language, concerned with the unstructured units that get computed by the computational system of human language, the lexical concepts are characterized by features, which are syntactically non-structured mental representations. But on the communication-theoretic view, how do we fix a concept? The relevant signals here - phonetic stimuli - will typically differ from speaker to speaker, and from dialect to dialect. Symbols which humans use to convey thoughts are rather abstract, contrary to animal communication - say the dogs' wagging of tails and their leaving scents on trees - where we have clear cues of a salient optic and chemical sort. The latter cases make it seem plausible that, eventually, signals can be matched

with interpretations; in the case of the symbols that humans use this is less clear.⁵

2.4 The Mystery of a Shared Language

The facts noted with regard to differences in phonetic stimuli and interindividual differences in conceptual understanding might provoke the thought that there *must* be some shared language in the sense of a “public code” which would relate “public pronunciations” and “public meanings”, or else we couldn’t communicate. The elements of this code will have properties that we all for some reason agree on. This is a rather abstract construct, more abstract than the I-languages mentioned above, which they presuppose and which are concrete structures internal to an organism.

But take the word *disinterested*, which most speakers use to speak of someone who is uninterested, although a few others take it to mean “unbiased”, or “unselfish”. In this concrete case, we somehow communicate in using the word, but the codes are different. Somehow we seem to be fortunate enough that human communication is mostly as efficient as dog communication. Nonetheless, who or what decides over what code one uses? One old answer appeals to “communicative intentions”, so that what I *intend* or *decide* to convey is what fixes what my words mean. But by bringing in the notion of intention - close as it is to the notion of meaning - we provoke more puzzles rather than less.

Efficient communication doesn’t seem to *depend* on an exact matching of interpretations. People partially converge in assigning interpretations to signals, but on the whole communication and understanding is a matter of more or less. As in the question of what it is to study the notion of “being near”, it seems unclear what it is to study “communication”, if this is to be a scientific account, not a matter of history or human affairs. Maybe the more assumptions Mary and Joe share the better they get along in using their respective generative linguistic faculties, and maybe it happens that the more they talk, the more they will converge on background assumptions. But it seems wrong-headed to stipulate that at some early point of such an interchange the people “do not (yet) communicate”, that is, that there is a minimal amount of a shared background that *has* to be there for a linguistic exchange to deserve its name. Whether it does or does not deserve that name seems like a (contextual) matter of human evaluation. If so, language is not, by its essence, social. Nothing *has* to be shared (although nobody

⁵To my knowledge there is little by way of theories of convention that would explain how interpretations are attached to signals. Even the simplest human concepts are of an extraordinary complexity, raising poverty of stimulus problems of the strongest form, considering also the rate of concept acquisition (see [Cho00], 63–66 for relevant comments). This suggests that concept acquisition is merely a labeling of concepts pre-existing in the mind. Labeling in this sense would be a brute-causal process, in which matters of rationality and conventional agreement play as little a role as in matters of the acquisition of syntax.

denies of course that language use is a social phenomenon). A speaker who says

(5) My dog is happy today

may be referring to a dog he owns, a dog he is making a walk with, a dog he is seeking to purchase, a dog he sketches on a canvas, or something else. Whether what he refers is “in fact” a dog, and whether it is *his*, is a matter of human value, much as is a judgement about whether the utterance was “appropriate”, given, say, that it was about a dog that is being painted on a canvas. Interpreters may want to figure out what exactly *my dog means* here, maybe with uniform results, maybe not. It is not some “shared expression” *my dog* (in a public code) which “generates a presupposition” (a phrase of Robert Stalnaker’s, to whom I will return). It is practical human reasoning that will decide, on an occasion of use and in ways that are independent of specifically linguistic principles, whether a speaker of (5) is taken to *own* a dog. Similar remarks apply to *happy* and *today*.

Thus we derive the preliminary conclusion that there is a non-contextual study of meaning in the form of a study of the language system’s SEM, but that the contextual aspects of the meanings of public speech acts - although no doubt they exist, as a matter of historical narrative, of norms and values - do not seem to lend themselves to a sort of scientific analysis that would be continuous with the study of SEM.

3 Content as Context

3.1 Dynamic Semantics

Semantics, the formal study of meaning, has mostly made language-world relations its topic, but has also tried to explicate linguistic meaning in terms of the relation of a sentence to the mind of a speaker who is using it. To some extent, dynamic semantics ([Chi95]), or rather some versions of it such as Kamp’s Discourse Representation Theory (DRT, [Kam81]), falls into the latter tradition. At the same time, it remains firmly rooted in a broadly Fregean approach to language. Thus “mind” is not primarily used as a psychological category here, and the focus has been rather on the model-theoretic study of inferential relations between objective propositional contents which are attributed to (and which help to individuate) epistemic states. But the idea of content is given a dynamic twist in that the semantic value of a sentence uttered in conversation is defined to be an update function that expands the body of shared presuppositions of the participants into a new such body. Thus we have a function,

$$f : C \rightarrow C,$$

where C is a context, a body of information or a set of presuppositions or propositional attitudes against which a new utterance is interpreted. Semantically speaking, C is a set of possible worlds, but this set is “presented” in terms of a syntactic *discourse representation structure* (DRS). It is something like a formula of predicate logic which is made sensitive to the way a hearer cognitively represent a certain amount of truth conditional information. DRSs are at the same time used as representations of the contents of the beliefs and other attitudes that a hearer forms in response to a discourse. In this model content and context are unified. A formal semantic representation of a discourse can be looked at as a *content* (a set of truth conditions) as well as a *context* (for the interpretation of further input). What a sentence is said to “do” is to update the semantic representation generated by the previous sentences, giving a new content as well as context as result. Naturally, this new perspective on meaning involves an increased involvement with pragmatics within semantics. A sentence meaning now itself consists in a dynamic potential for changing a context. It is not merely a semantic content, kept free of the pragmatic elements governing contextual changes.

3.2 Context as Common Ground

An empirical hypothesis under which the modelling of multi-speaker discourses in dynamic semantics proceeds, is that

“speaker and hearer engage in constructing a common ground, maybe from an already existing common ground constructed in earlier conversations” ([?], 25).

Over and above this anthropological hypothesis the meaning of an uttered sentence is typically regarded as

“an instruction to carry out a series of actions on a given database, with the end effect of incorporating the information given in the sentence into the database” (22).

Notwithstanding certain uses such a modelling decision may have, the question is what it tells us about human communication. Far from being a condition for the possibility of communication, the first above hypothesis restricts the resulting model to the possibly rare type of *common inquiry*, whose participants are unanimous in their goals and values. But agents do not exclusively or mainly strive for the truth, and for getting their findings as effectively as possible across to his fellows.

Even if they engage in common inquiry, agents do not let themselves be “instructed” by whatever is the content that speakers may program into a sentence. What does the meaning of a sentence used in an utterance have to do with the changes I perform on

my belief state? It goes without saying that an utterance is a natural object that may have causal effects on interlocutors. But agents in discourses are not generally speaking databases who could be brought to change by laws of cause and effect. How I change my mind is not a matter of my understanding of language, or my linguistic competence (alone). Human agents - and only those are of interest in the present chapter - don't change their mind by being coerced to do so. Understanding is a creative and deliberate process, not a causal one.⁶

It is also argued that for asserting, requesting, promising etc. to take place, interlocutors must attempt to recognize the speaker's intentions, and that this can only be if they cooperate. Speakers, as Grice put it, must

“make their contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which (they) are engaged” ([Gri75], 45).

Maybe humans tend to be cooperative as a matter of anthropological fact, but it is unclear whether linguistic theory should engage anthropology in this way. Also, there may not be such an “accepted purpose or direction of the talk exchange”, and although an agent may assume or come to the conclusion that there is, he may at any time become uncertain about whether unanimity really prevails. And even if he *does* act cooperatively, we may ask why. That he accepts the Gricean maxims is no answer to this question. We might try to construct an explanation that makes the interlocutor's action come out as an optimal decision in its own best interest, but this would be an explanation within the theory of rational decision. Whether the latter lends itself to explanatory rather than normative purposes is an open question.⁷

A hearer may not bark at some statement that a speaker puts forward, and he may have a good reasons not to do so. He may enter into a political agreement (a professor for example, who listens to the philosophical views of a freshman). Or it may not be useful to weigh the views of others, and to start an inquiry as to whether they are really correct. But it will be some such reasons, which may or may not be there, which explain the fact, if it is one, that a conversation steers towards a common ground, either

⁶Which also means: it is not a *mechanical* one, as in DRT, where discourse is looked at as a syntactic structure S in which sentences and sets of sentences (discourse segments) are distinguished. If it comes to the interpretation of the i 'th sentence S_i , a mechanical device called the *construction algorithm* operates on the syntactic analysis of the sentence and adds its phrasal constituents to the syntactic representation $K_i - 1$ of the discourse up to the sentence $S_i - 1$. The result is a syntactic representation K_i . Philosophically speaking, Fodor's computational theory of thought is operative here, which conceives thought processes as causal-computational relations between mental representations. Note that principled limitations to the Fodorian program argued for by Fodor himself in Fodor ([Fod00]) apply, if correct, in full force to the DRT program.

⁷Levi [Lev97] argues forcefully against the explanatory use of the theory of rational decision.

a political agreement or an actual shared agreement at the end of a common inquiry. In the same vein, we may observe that a *question* is something that an addressee tends to *answer*, although some agents do not. Or that agents take turns in conversation according to some order. But if we want to understand what is the driving force *behind* such phenomena structuring a discourse, explanations come to an end, and normative considerations prevail.

Generally speaking, we make sense of acts of uses of our linguistic competence in the very same way in which we make sense of other kinds of actions, such as the action of playing the piano, which involves use of a musical competence. Musical competence might be scientifically studied, or naturalistically. But a decision to do what we call “play the piano” will likely not, belonging as it does to the realm of human action. The theory of rational action may make this intelligible in some sense, but not necessarily in an explanatory one, as noted. If assumptions concerning common grounds, joint inquiries, cooperativity and so on are to enter linguistic explanation, large questions concerning the explanatory scope of a linguistic theory of this sort arise.

3.3 Changing the Context

Robert Stalnaker, one of the forefathers of dynamic semantics, argues that the “contents” of speech acts “should be understood in terms of the way they are intended to affect the situation in which they are performed” ([Sta98], 4-5). Now it may be that where communicative action is strategic, as in argumentative discourse, my choice of a sentence to make an utterance is a result of considering its potential effects. That however leaves open what this *choice* and its potential contextual effect has to do with the *meaning* that the sentence has in the language in which it is formulated. How speakers change their assumptions in response to an utterance is a process shot through with pragmatic and non-linguistic inferences exploiting background beliefs, social conventions, visual experience of the discourse context, and other things. While my *choice* of a sentence depends on features of the context, a *sentence* means what it does, in a non-contextual manner, by virtue of the rules of the grammar determining the language. A sentence like

(6) John said to Bill that he wants to wash him

is multiply ambiguous, and this is an essential fact about its meaning. That has to do with the Binding Theory, not with how speakers intend to change contexts.

In short, to the same extent that Stalnaker’s “contents” depend on how speakers change contexts through their utterances, they are of dubious value for an analysis of the notion of linguistic meaning. Stalnaker may be right that “[i]t is simply an obvious fact that an assertion changes the context” ([Sta98], 12), but there is no way to compute, from

the sentence used in the assertion, the solution to the question of how the context will change. So what explains the causal effects ensuing from using language?

3.4 Discourse as Inquiry

A context in dynamic semantics usually is an *epistemic* context, modeled by a set of assumptions or beliefs, which in turn are analysed as a set of possible worlds in which they are true. If meaning is a context change, it would thus be a matter of the alteration of our propositional attitudes (beliefs, presuppositions). These alterations (context changes) are acts specified under intentional description. That in itself doesn't make the enterprise a normative one. Nonetheless, normative aspects surfaced above for other reasons, and in general, it is hard to see how determining what beliefs a person holds in a context could be an entirely empirical matter. In figuring out beliefs we make assumptions concerning which beliefs it would be *rational* for the person to have in the given context, and that question will guide us to figure out which ones she actually has. Similar considerations apply to the study of epistemic context change, which seems inseparable from the question of how they should change. Hence normativity enters, and meaning, if explicated in terms of the notion of context change, will be to some extent normative, too.

Normativity may quite clearly enter even in the way that considerations about what someone would rationally believe in a context may overrule empirical data about what he says. Somebody who sincerely asserts "I am going downtown" may be interpreted to mean "I am going uptown" if it simply makes no sense at all to interpret him to think or believe that he goes downtown. But note that this shows nothing about the *meanings* of these sentences, which is just as it should be: the meaning is what it is, by the rules of language, and it doesn't change in the light of considerations of rationality. But context changes do depend on such factors: it doesn't have to do with meaning (alone).

To give an expression to the normativity of epistemic state changes in discourse, we might say that characterizing agents in discourses in terms of their propositional attitudes, as dynamic semantics does, is really to describe them as undertaking certain attitudinal *commitments*. A basis for this consideration might be that forming an attitude is what makes an agent committed to make judgements that correspond to this type of attitude [Lev97]. For example, if an agent comes to the conclusion that a proposition is fully (i.e., non-conditionally, non-partially) true, he undertakes a rational commitment to a disposition to assent to it as well as its logical consequences when prompted to do so. The shift to talking about beliefs as attitudinal commitments helps to bring out just why discourse updating and context change must be a deliberational matter. Sometimes it may well be the case that you change your commitments routinely in the light of new information. But often you will have to weigh risks against potential gains in changing them: you may be deliberating on whether to trust a person whom who previously

distrusted, and while trust may be the option that you would like most and that has the highest value, it might be the most insecure.

Judgements corresponding to the various propositional attitudes allow us to make distinctions between propositions on any chosen level of fine-grainedness. We might distinguish

- propositions that are judged fully true, forming a background for reasoning and inquiry that is not now questioned;
- propositions that are, relative to this background, judged uncertain or probable, where the probability is subjective, and
- propositions that are judged valuable, in accordance with one or the other system of value commitments.

The first propositions I will call (full) *beliefs*, the second *potential beliefs* or *states*, and the third *outcomes*. To indicate how a normative model of discourse updating might get some structure, we might identify the commitment (to the truth of) a proposition with the set of the commitments generated by it, i.e. its (deductive) consequences. This gives us a notion of the 'content' of an attitude, but, of course, a strictly normative, not a semantic one.⁸

On this picture, it is not reasonable any more to inquire into the “nature” of propositions. The term proposition as just conceived is not a natural object but a technical notion that has exactly the meaning we give it. What meaning we give it will depend on our purposes: what propositions we wish to distinguish, and how we distinguish them. Independently of given interests and concerns that define what the propositions are that are the objects of a deliberation at a moment of inquiry, and thus when two of them are identical, the question of propositional identity is not meaningful. It cannot be settled by appeal to some absolute authority whether “Hesperus is Hesperus” and “Hesperus is Phosphorus” express the “same proposition”. It will depend on the contexts in which we introduce and use these notions that the question will be decided in one or the other way, relative to given purposes and concerns. Language doesn't tell us whether to count them as the same or not, or whether to identify the messages conveyed by “Paul drove Paula home” and “Paula was driven home by Paul”.

A major reason why contextuality is massive within rational inquiry is that there may not be general rules for partitioning the space of propositions, or for setting up an ultimate partition of best options that one chooses to consider in the course of one's current decision problem. In a similar way, there seem to be no general rules for determining questions in discourse and for abducting a set of potential answers for them. That is

⁸Levi [Lev91] adopts such a notion.

to say: question answering in discourse is not a computational process, but a rational one. There seems to be no clear sense in which the statement that a question is a set of potential answers could be a statement within a descriptive linguistic theory. Not even the construct of a shared language that I argued against above can make the deliberative and normative aspects in constructing and answering a question go away: language doesn't have the effect to make you construct, by itself, a set of potential answers relevant to the speech situation. Consider the interrogative

(7) Do you accept Credit Cards?

spoken by me when calling a restaurant. It does not come with a partition of potential answers, nor does it determine, by itself, what the question is in the first place. It may not in fact be a question, but, for example, an overt or hidden threat. If it is a question, it may be inferred that I want to know whether one of the accepted Credit Cards is the one I own; whether I can pay with one of my Credit Cards, which are most probably among those accepted in a restaurant; whether I can patronize this restaurant; etc. Having inferred the question, a Hearer abducts a set of strongest consistent, exhaustive, and pairwise inconsistent potential answers. The final outcome of this abductive phase of deliberation contains those potential answers that are the options between which one has decided to restrict one's choice. It is exactly as fine-grained and as coarse-grained as one finds it useful. For all it seems, no general rules for forming a partition can be given. In a similar way, it seems that no general rules can be given for what demands of information an agent should have, or with what degree of caution he should reject hypotheses.

In sum, what I have tried to bring out is that what dynamic semantics really seems inspired by is a vision of *inquiry*, not meaning, and that eventually that vision will have to show its true colours as a *normative* one. We might indeed define a technical notion of "meaning" so as to consist in the set-difference of the Hearer's partition before and after the utterance of a sentence has taken place. In this way we would read off the meaning of a sentence from the inductive inference (an elimination of options) that an agent performs. That this is a non-linguistic notion of meaning seems rather clear, however, and it might be more useful to call it a notion of *significance*. In any case, its *massive contextuality* – the fact that it will depend on contextually given human interests and values in an almost arbitrary manner – makes it an unlikely candidate for a naturalistic inquiry into human nature and language.

4 Stalnaker's Critique of Dynamic Semantics

Dynamic semantics is often said to have one of its origins in Stalnaker's work on assertion. Nonetheless Stalnaker ([Sta98]) criticizes it precisely for taking the pragmatics

of contextual change to define *semantic* content. The attack recommends separating a purely semantic core notion of meaning from pragmatic aspects of meaning having to do with the use of expressions that have that semantic content. Linguistic phenomena that supposedly motivate the introduction of dynamic semantic values (functions from contexts to contexts), Stalnaker argues, can be accounted for just as well with classical models of contexts and propositions (sets of possible worlds). It is the *point* of a meaningful speech act to change a context; but this is not what defines the *meaning* of a sentence.

All this is intendedly Gricean in spirit: Grice also reconciled the traditional non-contextual semantic analyses of logical expressions by adding a non-semantic account of conversational maxims to them. These would leave the assigned semantic contents untouched but lead us to the specific interpretations that a context might require. But in the case of Stalnaker we may ask exactly what we may ask in the case of Grice: whether the notion of semantic content that we protect by adding pragmatics to the theory of language is an artificial phenomenon rather than a natural kind.⁹ Both the notion of proposition and the notion of meaning just used are fictions if the above discussion is right: they are artificial abstractions from the massively contextual interplay of a myriad of faculties as they enter public discourse. We will now try to substantiate this prior conclusion above in the light of Stalnaker’s account.

It is essential to Stalnaker’s argument that the proposition, in his sense, that is asserted (or “proposed for acceptance”, as Stalnaker says) is fixed before we look at the contextual dynamics that the assertion of the sentence expressing that proposition affects. Specifically, if a context changes through the assertion of a sentence, it does so in two ways: first the set of mutually shared assumptions adjusts to the fact that a *particular sentence with a particular content* has been asserted (first context change); secondly, the proposition proposed for acceptance is either accepted or rejected (second context change) ([Sta98], 8). Semantics is dealt with when the first change has happened. Pragmatic principles control the second. Both changes are formally modelled as modifications of sets of possible worlds modelling the information that is compatible with the shared assumptions of a group of speakers at a particular moment of their discourse.

Note that Stalnaker’s second step of contextual change involves adopting or not adopting a belief that will in general not be a logical consequence of what is believed already. In this sense the inference that is drawn to its truth is ampliative or inductive. It will necessarily carry risk in the sense alluded to above. An inference that proceeds

⁹As is sometimes said: Semantics was saved by inventing pragmatics. But what if there is no semantics? It is not clear whether there is not, but it seems necessary that with respect to natural languages the existence of a level of semantics and “propositional contents” is an empirical matter entirely. See e.g. Hornstein [Hor83] for general arguments that there is no such level, Jennings [?] for empirical refutations of Gricean stipulations of logical constants in natural language, and Dudman [Dud91] for the specific case of English *if*.

from a body of beliefs that are held true to a new such body involves a risk of error from the viewpoint of what is currently held true. While it is trivial to model such an induction as a change in the set of possible worlds capturing what is taken for granted at a stage in discourse, what is decisive in a theory is the reasons why such changes happen. One cannot build the rational structure and the human values that enter into the reasons for such changes into a number of “conversational maxims” in Grice’s manner, or, as Stalnaker’s expresses this, a “number of truisms about conversation as a rational activity” ([Sta98], p. 18).

Let us look at the fabrication of Stalnaker-propositions in some detail. Recall Stalnaker’s first type of context change, caused by the “manifestly observable event” ([Sta98], 8) that a “statement was made”. But together with it and “standing assumptions” that are part of the prior context we can “infer, (...) not only that the speaker uttered certain sounds, but also that she uttered an English sentence, and that she is saying something to us” (8–9). Further assumptions (cooperativity, competence), Stalnaker goes on to assume, makes inferable from observational data, in fact, that we have a token of *an English sentence type whose meaning and presuppositions in the event of its utterance we know*. For example, when Phoebe asserts

(8) I can’t come to the meeting – I have to pick up my cat at the veterinarian,

the first context change that takes place according to Stalnaker is that we now presuppose that Phoebe owns a cat, and that this is now a shared presupposition.

But why do we do such a thing? Stalnaker answers: because after we witness the utterance of (8), we presuppose, not only that it took place, but also that Phoebe “is making whatever presuppositions are required to make her utterance intelligible and appropriate” (9). This in turn is because we have the prior belief that Phoebe is a “cooperate and competent participant in our conversational enterprise” (9). From these assumptions, together with the further one that the expression *my cat* has as part of its “appropriateness conditions” that its speaker owns a cat, we presuppose that Phoebe means to presuppose that she owns a cat, and presumes this information to be shared. Finally, “since it is presupposed that presuppositions are shared information, he will accommodate by presupposing it himself” (10).

It should be clear that all of this takes a lot of good will. In fact it seems that if we ask again why we do any of these things, what we are pointed to is the fact that there are certain “standing assumptions”, both about what Phoebe is like, and what it is for certain expressions to be used “appropriately” (note how the shared language assumptions is presupposed here). From this it would seem that whatever is needed to get us to draw all those inferences is coded into our premises. Be that as it may, it remains that all of the inferences and conclusions mentioned that we draw from an utterance event are *inductive* ones in the sense above: but that by itself implies that we

need not draw them. It is not enough to point out that we will “typically” draw them, or that there is a “disposition” to do so. Drawing an inductive inference is to go beyond present commitments, and this is and remains a deliberative matter.

To sum up this discussion, it seems that at the point where we have got a Stalnaker-proposition (an absolute, non-contextual item of information, a set of possible worlds) and a presupposition associated with it, massive inferencing has already taken place. When it comes to Stalnaker’s second kind of context change – where a proposed proposition is accepted or rejected – what we have is a product of this second kind of context change (a belief change, an either routine or deliberate expansion, in the light of an utterance event) already. The determination of the proposition proposed for acceptance in a speech act itself requires making reference to pragmatic and inductive principles for changing one’s context or doxastic dispositions. This is evident even given the assumptions Stalnaker makes in his example. It becomes abundantly clear if the shared language assumption fails, or the inductively inferred cooperativity and competence assumptions are given up. Semantic content in Stalnaker’s sense is itself largely a matter of pragmatics and inductive inference, and not given independently by “language” itself. Stalnaker himself in fact remarks that we may happen not to be “willing” to infer that Phoebe has a cat. What is surprising is that a discourse should have to be rated “defective” ([Sta98], 10) if this is so (it would appear that such clearly normative considerations have no place in semantic theory). Due to the defectiveness, Stalnaker argues, some “backtracking and repair” is required. We may conclude from this that even for Stalnaker, in the *general* case how a context changes in the first sense is a matter of pragmatics and reasoning, hence is itself a context change in his second sense. It seems then that Stalnaker has no point in claiming that dynamic semantics mixes pragmatic matters into its notion of a semantic value. Stalnaker himself doesn’t isolate a clearly semantic notion of linguistic content that would be logically prior to the context changes that the dynamic semanticist talks about.

Let us finally apply these conclusions to a concrete and famous example of Partee’s by virtue of which Stalnaker aims to show that dynamic semantic values aren’t needed. Imagine the following well-known pair of examples, said in the same initial context:

(9a) Exactly one of the ten balls is not in the bag.

(9b) Exactly nine of the ten balls are in the bag.

One argument for dynamic semantic values went as follows: If the contents of (9a) and (9b) are modelled as sets of possible worlds, their contents will come out as exactly the same. This, dynamic semantics has claimed, is a problem since if we continue a discourse started with either sentence with the new sentence

(10) It is under the sofa

this sentence will be interpreted differently in the two cases. The possible world analysis doesn't predict this.

The supposed problem for the possible worlds analysis vanishes, Stalnaker argues, if his first kind of context changes is taken seriously. For since different words were used in the utterances of the respective sentences (9a) and (9b) to get across what is in some technical sense the "same content", the context will change in different ways. An utterance of (9a) makes an antecedent for *it* in (10) uniquely salient, and (9b) fails to do so. This is a fact that "can be presumed to be available" to the speakers, hence, if sets of possible worlds model the context of interpretation prior to the utterance of (10), these sets will be different in the case of (9a) and (9b).

While that may be right, it leaves open, as Stalnaker acknowledges, *why* they are different. But this is what we would like to explain. The set difference is no more than an expression of what we wish to understand. And there is a problem of understanding, for, as noted, it simply won't do to say that a sentence is a context changer. Sentences, as noted, do not do such things, and even if *it* in (10) picks up the one ball made salient through an utterance of (9a), it is not facts about (9a) qua English sentence which cause this. Neither does it help to stipulate that "the point of a speech act (...) is to change the context" (1998, 8). Stated as such this cannot be right, as it can only be the point of an assertion to change the context *in a certain way*. But if thus understood, we derive what is to be explained by introducing a new entity, a communicative intention, which fixes the point of an assertion. But the phenomenon of the specific context changes we have to explain *consists in* differences relating to making a certain communicative act, and we cannot appeal to the latter to explain the former.

Maybe we could call the process that leads to assign the missing ball to *it* a default inference. If so, a default inference is one that needs to be *drawn* by a hearer, and it is again not clear whether this is a matter of linguistic form alone. The assignment of a referent to *it* in the case of an utterance of (9b) - namely, an assignment of either the bag or the missing ball (which might as well have been made salient by mentioning the ones that are *there*) - is in either case clearly not an inference based on linguistic form. Of course there is nothing wrong in inferences not based on linguistic form, but an explanation won't benefit if facts about what inferences we may or may not draw from linguistic facts are packed into our notion of linguistic meaning.

Even if the assignment of the missing ball to *it* in the case of an utterance of (9a) were a linguistic matter, a postulation of richer semantic representations than is offered by Stalnaker's austere apparatus of possible worlds might be very useful. Thus Kamp's DRT suggests formal semantic representations for (9a) and (9b) which differ in that one contains a discourse referent which subsequent pronominal discourse may pick up, while the other does not. Here the existence of this difference in formal structure is a way of *explaining* the set difference that Stalnaker mere takes as given. Furthermore, DRT offers a specific construction algorithm operating on the syntactic structures of the sentences

in (9) which yields the difference in formal semantic representation. In other words, the formal difference in representation and content is traced to linguistic mechanisms, not invited inferencing. This is good because it doesn't let semantic contents be made up of non-linguistic facts such as invited pragmatic inferences. It is bad if I am right that the inferences leading to an assignment of a referent to *it* both in the case of an utterance of (9a) and (9b) are of a non-linguistic sort, in which case a non-linguistic inductive inference would be wrongly modelled as a linguistic mechanism.

I am not insisting on a point that boils down to the question how we define language. The question how rich a notion of language we can afford if we are after an explanatory theory of language. Suppose it was right that

“speech acts are context-dependent: their contents (and so the way they are intended to affect the situation) depend not only on the syntactic and semantic properties of the types of the expressions used, but also on facts about the situation in which the expressions are used” ([Sta98], 5).

Then it would be true that computations that the mind performs when constructing a semantic representation of an input sentence cannot be *local* in a sense suggested to be necessary for a formal account of human knowledge of language in the beginning of this paper. If the notion of content employed in the quote was adopted, principles for computing a semantic representation couldn't be defined locally over features represented in the human head, let alone an encapsulated language module. If the non-linguistic environment in which expressions are used relevantly enters the computation, it will be dubious whether we can still be talking about a theory of semantic representation as part of a computational theory of cognition.¹⁰

I finally and briefly mention another problem in the Stalnakerian picture, relating to questions of methodology. The strategy is “to get clear about what language is for – what it is supposed to do - before explaining how it does it” ([Sta98], 4). But language is a natural phenomenon, and it seems in general dubious to ask for a natural phenomenon “what it is for”. Language is used to change contexts for sure, but the question is whether this teleological talk is more than a matter of historical narrative. Communicative purposes may be invoked in informal description of the biological phenomenon, but abandoned as we progress to explanatory theories, in which we discard notions of function, intention, and purpose. The concrete “mechanisms and devices that particular languages may provide for doing what is done in a discourse” appear to be what we ultimately aim for, and I know of no reason to believe that these mechanisms

¹⁰See Fodor [Fod00] for a systematic elaboration on the locality properties needed for cognitive processes to figure in a computational theory. Note again that if dynamic theories of discourse took Fodor's point seriously, they would have to exclude abductive processes – which Fodor argues are non-local – from their account of semantic representation.

will have to vindicate our common sense intentional description of “doing what is done in a discourse”, or even be understood as means to achieve such doings. There remains a question, then, about the status of the “descriptive apparatus” that Stalnaker proposes to develop to study the interaction of “content and context”. If we study a natural phenomenon, maybe we should aim not for some such “apparatus”, but an explanatory framework within which we can try to find real categories of human cognition.

5 Contexts in LFs

5.1 Context Variables

The notion of meaning employed in section 2 above was a non-contextual one, leaving matters of interpretation relating to the utterances of expressions aside. I granted that in the interpretation of utterances contextual factors of any imaginable sort are likely to enter the process of interpretation. I will now ask whether there are linguistic ones among those factors.

Assuming it is not words which refer to things but people which do that on occasions of using them, we can start by noting a familiar contextuality with respect to indexicals like *I* and demonstratives like *this*. One says an expression like *I* is “context-dependent”, meaning the referent shifts on who uses the expression. In the case of *I*, the very act of using the word is sufficient to fix the referent. But that should be said to be a feature of the word *I* relating to *non-contextual* features of its meaning, for what is contextual is the referent determined by features of its meaning in an act of using it.

Similarly, what is non-contextual about the word *this* is that it must be accompanied by an act of reference - a speaker must use it to refer to something or other - even though only the context of use will tell what object is being referred to. A pronoun like *she* would appear to be intermediate between these extremes, constraining on linguistic grounds the referent to be a person that is female, but leaving the choice of actual reference within those confines to the speaker.

In line with these observations concerning the linguistic and the pragmatic (or use-theoretic) aspects of an expression which enter interpretation, Higginbotham [Hig88] suggests stating truth conditions for context-dependent utterances like

(10) She is lazy

in the following manner.

(11) If x is referred to by *she* in the course of an utterance of (10), and x is female, then that utterance is true just in case $\text{lazy}(x)$.

As Higginbotham notes (p. 34), this makes the truth apply to utterances in an 'absolute' fashion: there is no notion of truth-in-a-context in (11). Contextual features that influence the act of interpretation are mentioned in the conditional's antecedent, and the right hand side of the biconditional in the conditional's consequent registers their effects on the interpretation.

The semantics in the style of (11) needs further refinement for cases in which a speaker doesn't use a noun phrase to refer to a thing but merely restricts the range of things that the phrase may refer to. Examples are (12), on the reading where *he* is bound by *everyone*, and the famous (13):

(12) Everyone likes the book he read

(13) The murderer (whoever he is) is insane.

Even though it is required by the linguistic form of (12) alone that everyone read a single object and that this object was a book, otherwise the referent is indeterminate. Inspection of the context of an utterance of (12) will provide the referent(s). Similarly in (13), where no particular person is referred to. Rather, the noun *murderer* is used to restrict the quantification to persons who murdered someone, and the speaker's role is, in Higginbotham's phrase, to *confine the range of the noun to things such that* certain contextually given conditions are fulfilled. Symbolizing these conditions by a free, second-order variable X serving to encode contextual restrictions on the range of quantification of expressions with generality impact, the normal form of stating truth conditions for something like (13) becomes:

(14) In any utterance of (13) in which the speaker confines the range of *murderer* to things x such that Xx , that utterance is true just in case: $\text{insane}((\text{the } x)(\text{murderer}(x) \text{ and } Xx))$

Context variables will have to be employed in full generality, as it is standardly the case that ranges of quantification of general expression are usually not their full ranges but appropriated to a context. In capturing contextual effects on interpretation in this manner we are not delving into a study of the multifarious facets of communicative contexts, but specify what is required, on the linguistic side of the interface with the pragmatic systems, to determine the content of an utterance.

A significant virtue of context variables is manifest in the semantics of names. As Uriagereka [Uri98] notes in the context of other examples, there is nothing wrong, say, in accepting that while Pizza Quattro Stagioni tastes great in my local pizza place, it doesn't taste great today (considering the instance here in front of me). How is that consistent (it obviously is)? Apparently, by noting that human languages allow us to distinguish between individuals (such as pizzas) *at* certain events in which they

participate (such as the current eating event), on the one hand, and self-standing or non-contextual individuals, on the other. Obviously, we are ill-advised to follow traditional practice and formalize names as simple constants denoting or labeling non-contextual objects in the world. It may be essential to the objects as referred to by us that they occur in a certain context in which they play a certain role. The pizza Quattro Stagioni in today's dining event is a different Pizza Quattro Stagioni from the one I know and love. It seems like a crucial fact of human cognition that it is in contexts that speakers present notions in specific respects.¹¹ In the next section I present observations of Raposo and Uriagereka [RU95] that, furthermore, it is in contexts that speakers introduce *other* contexts, hence that context variables take scope over one other, much as quantifiers do.

5.2 Relative Scope of Context Variables

The way context variables are *ordered* with respect to each other appears to account for differences in predication that can be observed in examples like (15):

(15a) The city has poor neighborhoods

(15b) The city is (one) of poor neighborhoods.

There is an obvious semantic difference here, relating to the fact that in (16a), we consider the city in the context of a particular issue, namely poor neighborhoods. With respect to that, the city has those neighborhoods, but nothing is said as to whether it has rich ones as well. That is not so in (16b), where the poor neighborhoods are what we might call a “standing characteristic” of the city, which holds of it in a decontextualized fashion. In general terms, while in (15a), the context of the city is grounded on that of the predicate, it is the other way round in (15b). The ambiguity can be observed generally, as in

(16) Marlowe loved women

which can be interpreted as a statement about a standing characteristic of Marlowe, a factual feature of his personality, but equally as a report on an erotic event or a sequence

¹¹From this point of view there is a clear disagreement with Stanley's claim ([Sta00], 400) that overt expressions that are neither indexicals nor demonstratives are “not context-dependent”. If what I said is right, there is a sense in which we would rather have to say that literally every expression, no matter its category, is “context-dependent”. There is no fixation of reference without, for example, taking speaker's non-linguistic intentions and presuppositions into account, even in simplest cases. On the other hand, given the above context variables one sees the affinity to Stanley's approach in his formulation that “any contextual effect on truth-conditions that is not traceable to an indexical, pronoun, or demonstrative (...) must be traceable to a structural position occupied by a variable” ([Sta00], 401).

of such events in Marlowe’s life. In the latter case, we might say that it’s the loving of women that we talk about, and which figures as the sentence’s topic.

Such observations suggest that assuming every quantificational expression comes with a context, and assuming further that each predicate comes with an extra event argument [Hig83], each simple “*S* is *P*” predication leaves it open which context is set into which. Either the predicate is topicalized (moved to a site marked for topicality at SEM), in which case we talk of some event in the context of which we consider some particular thing as its subject. Or the subject is topicalized, in which case we introduce the predicate in the context of that fully individualized object, of which it will now hold as a standing characteristic. A striking illustration of Raposo and Uriagereka [RU95], who suggest this idea, is (17) and (18):

(17) I consider the sea the frog green

(18) I saw the sea the frog green.

There is a canonicity effect with respect to the greenness of the sea and the frog, respectively, in (17): the sea green will very possibly differ from the frog green (the former being, say, more blueish). That typicality of the colour is not invoked in the case of (18), where as a result of the light and visual circumstances we may well have seen both the sea and the frog in exactly the same green.¹² We can explain these observations in the light of the prior considerations. If the green is predicated in the context set up by the subjects *sea* and *frog* considered as topics, it is the green which is canonical to these items. If, on the other hand, what we talk about is an event of seeing greenness, and the sea and frog are as it were reduced to participants in this event, the green that we see may be whatever it is: the sea and the frog will have *it*, and it will not in any essential way be the sea’s green or the frog’s green.

In consequence, there is evidence that it is not merely that context variables are instantiated in context on the basis of whatever information and cues are contextually given, they are also instantiated sequentially one *after* the other in the process of constructing semantic representations.

6 Conclusions

The emerging picture appears to be this. I have started arguing that a theoretically promising notion of language must be wary against the contextuality that is relevant and pervasive to the study of human communication and purposive language use. While it is true that we use language to communicative purposes and the transmission of thoughts,

¹²Data about auxiliary selection in the Spanish counterparts to these examples make the difference of predication overtly manifest in the syntactic form.

this need not explain us why languages generate structures and meanings in the way they do. Linguistics on this picture is a science of expressions, not a study of actions specified under intentional description, hence completely remote from the normative issues that unavoidably arise in the latter.

This specifies the role of the non-linguistic context given by the uses to which expressions are put. Semantics, if conceived as the study of rules by which linguistic sounds relate to meanings, is non-contextual. On occasions of use, matters of linguistic form are contextualized, but the massive contextuality found in inductive inferencing does not seem to lead to a foundational clear and fruitful notion of content. On the other hand, none of this hinders variables encoding context to enter linguistic description and to influence our conception of LF, and it even may give rise to reconsidering certain traditional conceptions in the theory of object and event reference.

References

- [Chi95] G. Chierchia. *Dynamics of Meaning*. University of Chicago Press, Chicago, 1995.
- [Cho95] N. Chomsky. *The Minimalist Program*. The MIT Press, Cam., Mass., 1995.
- [Cho00] N. Chomsky. *New Horizons in the Study of Language and Mind*. Cambridge University Press, 2000.
- [Cla92] H. Clark. *Arenas of Language Use*. University of Chicago Press, Chicago, 1992.
- [DiS99] A.M. DiSciullo. Formal Context and Morphological Analysis. In *Modeling and Using Context*. 1999.
- [Dud91] V. Dudman. Interpretations of 'if'-sentences. In F. Jackson, editor, *Conditionals*, pages 202–232. Oxford University Press, 1991.
- [Fod98] J. Fodor. *Concepts*. Clarendon, Oxford, 1998.
- [Fod00] J. Fodor. *The mind doesn't work that way*. The MIT Press, Cam., Mass., 2000.
- [Gri75] P. Grice. Logic and Conversation. In Morgan J.L. Cole P., editor, *Syntax and Semantics*, volume 3: Speech Acts. Academic Press, New York, 1975.
- [Hig83] J. Higginbotham. The Logic of Perceptual Reports. In *Occasional Papers from the Center for Cognitive Science*. MIT, 1983.
- [Hig88] J. Higginbotham. Contexts, models, and meanings. In R. Kempson, editor, *Mental Representations*, pages 29–48. Cambridge UP, 1988.

- [Hin01] W. Hinzen. The pragmatics of inferential content. *Synthese*, 128(1–2):157–181, 2001.
- [Hor83] N. Hornstein. *Logic as Grammar*. MIT Press, Cam., Mass., 1983.
- [Hua95] J. Huang. Logical Form. In G. Webelhuth, editor, *Government and Binding Theory and the Minimalist Program*, pages 125–176. Blackwell, Oxford, 1995.
- [Jen95] R. Jennings. *The Genealogy of Disjunction*. Oxford University Press, Oxford, 1995.
- [Kam81] H. Kamp. A Theory of Truth and Semantic Representation², booktitle =. 1981.
- [KR93] H. Kamp and U. Reyle. *From Discourse to Logic*. Kluwer, 1993.
- [Lev91] I. Levi. *The Fixation of Belief and its Undoing*. Cambridge University Press, Cambridge, 1991.
- [Lev97] I. Levi. *The Covenant of Reason*. Cambridge University Press, Cambridge, 1997.
- [Nun95] G. Nunberg. Transfers of Meaning. *Journal of Semantics*, 1, Winter, 1995.
- [RU95] E. Raposo and J. Uriagereka. Two Types of Small Clauses (Toward a Syntax of Theme/Rheme Relations). In A. Cardinaletti and M. T. Guasti, editors, *Syntax and Semantics*, pages 179–206. Academic Press, 1995.
- [Sta78] R. Stalnaker. Assertion. *Syntax and Semantics*, 9: Pragmatics, 1978.
- [Sta98] R. Stalnaker. On the Representation of Context. *Journal of Logic, Language and Information*, 7:3–19, 1998.
- [Sta00] J. Stanley. Context and Logical Form. *Linguistics and Philosophy*, 23:391–434, 2000.
- [Uri98] J. Uriagereka. A Note on Rigidity. *Linguistics Today*, 22: Possessors, Predicates and Movement in the Determiner Phrase, 1998.
- [ZS92] H. Zeevat and R. Scha. Integrating Pragmatics into Update Semantics. In Stock O. Ortony A., Slack J., editor, *Communication from an Artificial Intelligence Perspective*, pages 17–34. 1992.