Context and Interpretation

Erich H. Rast

Institute for the Philosophy of Language (IFL)

Universidade Nova de Lisboa

erich@snafu.de

Abstract

In this article, three shortcomings of parameter-based contextualist and relativist accounts of linguistic context dependence are laid out and then an alternative approach is sketched in which context is based on the assumptions of individual agents. First, context does not determine semantic values, but discourse participants do so when they interpret an utterance. Second, a fruitful notion of context must account for more than the mere variability of extensions from assessor to assessor. And finally, the standard two-dimensional apparatus has technical limitations that make it unsuitable for expressing fine-grained models of contexts. Instead of modeling contexts as complex parameters, contexts should, following Stalnaker, be regarded as sets of interpretative assumptions. Interpretation can then be taken as an inferential process on the basis of such assumptions. In a given interpretation situation, these assumptions may be generated by revising an agent's beliefs by his second-order beliefs about what the speaker believes.

1. Introduction

A variety of contextualist positions about the content conveyed by an utterance (or, conveyed by a speaker when uttering a sentence) have been defended in the past—see for example Recanatí (2004), Stanley and Szabó (2000), Stanley (2002), and Travis (2008). Recently, the trend has shifted towards assessment relativism—see for example MacFarlane (2005, 2007ab) and various contributions in Carcía-Carpintero (2008). In this article, I argue that a third position about context

dependence in communication, which may be called the interpretation model of context (IP), is generally more adequate to represent and explicate linguistic context dependence. This position has rarely been taken explicitly into consideration in course of recent controversies about linguistic context dependence, even though it has a longstanding tradition in AI research. According to IP, literal semantic content is interpreted by individual agents on the basis of their varying assumptions. My arguments against contextualism, henceforth abbreviated CT, and relativism, henceforth abbreviated RT, do however only apply to versions of these positions that implicitly or explicitly rely on two-dimensional modal logics in which contexts are considered parameters. This assumption is not made in all versions of CT. For instance, Recanatí (2004) does not seem to advocate it, and IP is itself also a form of moderate contextualism.

My critique on parameter-based CT and RT is not meant to be a general critique on regarding contexts as parameters; it is only a critique on to the way in which context dependences are resolved by context when a typical two-dimensional logic like that of Kaplan (1989) is assumed to be working under the hood. As I will argue, not even the saturation of what Perry (1997) has called automatic indexicals such as 'I' and 'now' can be adequately explained in the way stipulated by Kaplan's Logic of Demonstratives and similar approaches.

The remainder of this article is structured as follows. First, I will briefly lay out parameter-based versions of CT and RT and briefly discuss their differences. Then I will lay out the descriptive and explanatory shortcomings of such two-dimensional approaches, and thereafter I will sketch the alternative view of regarding context as a set of agent-dependent assumptions.

2. The Main Positions

2.1 Parameter-based Contextualism

In a moderate form of CT with two-dimensional semantics under the hood, the linguistic meaning (character) of a sentence is evaluated with respect to a context parameter. The result of this

operation is the semantic content of the sentence in the context, which is then evaluated with respect to some circumstances of evaluation (CE; another parameter) and yields an extension. Initially, relevant features of the CEs such as time and world are identical with those in the context, i.e. the time and world of the context, and only modal operators for tenses and doxastic or metaphysical modality change these features of the CE. As a typical example, the literal meaning of 'I was hungry' in a context *c* yields a content that with respect to a CE *i* is true if there is a time interval before the time of *i* (= the time of *c*) during which the speaker of *c* was hungry, false otherwise. As Kaplan (1989) has laid out in detail, the central idea behind this mechanism is that the context provides the referent of an indexical and this object forms part of the semantic content of an utterance. Tense and modal operators in turn may change the extension of a given content by implicitly quantifying over CEs, but they do not affect the way in which indexicals are saturated on the basis of their linguistic meaning (character). Such a two-dimensional semantic approach accounts for the rigidity of indexicals, as these are evaluated independently of ordinary modal operators in the logical language, and is at least to some extent compatible with the direct reference view while in the meantime encoding the rich descriptive meaning of indexicals as their character.

2.2 Parameter-based Relativism

Following Stojanovic and Predelli (2008), a generic version of RT can be obtained by loosening the initial coupling between contexts and CEs and enriching CEs with additional features such as a time and place of assessment or the assessor itself. The same semantic content obtained from the same linguistic meaning with respect to the same context parameter can then be evaluated with respect to different CEs. As an example taken from Lasersohn (2005), 'Roller coasters are fun' in a context *c* yields a content that with respect to a CE *i* is true if roller coasters are fun for the assessor of *i* at the time of *i,* where the time of *i* = the time of *c*.[1] Since roller coasters may be fun for one assessor

1 MacFarlane's temporal relativism is not assumed in this example. Notice that MacFarlane (2007c) doesn't agree with Lasersohn on faultless disagreement in case of predicates of personal taste; cf. Stojanovic (2008).

while not being fun for another assessor, and there is nevertheless a strong feeling that the two assessors talk about the same, the relativist asserts that different assessors may disagree about the same semantic content and both be right in the sense that this content may come out true from the point of view of the assessor of *i* and false with respect to a CE *j* with another assessor. This also allows for 'roller coasters are fun' to come out true with respect to the context, CE pair <*c, i*> and 'roller coasters are not fun' to come out true with respect to <*c, j*>. So the discourse participants can make seemingly contradictory statements that are both true relative to their personal vantage points. It is obvious that CT can be enriched to make similar predictions as RT by putting the respective assessors or standards of assessment into the context parameter.

There is, however, still a difference between such versions of CT and genuine RT. When traditional contextualism is extended by adding assessors to the context parameter the corresponding disagreement between two discourse participants is a disagreement about different semantic content. Relativists like MacFarlane and Lasersohn have insisted, on the basis of varying sets of examples, that this is counter-intuitive. As MacFarlane (2008) lays out, when somebody asserts something about the future an utterance might be underspecified *now* even though it turns out true or false *later* from the perspective of the indeterminacy of the future view defended by MacFarlane. Nevertheless, when someone for example utters 'On 19. 4. 2010 it will rain in Paris' on 14. 2. 2008 and someone else utters 'On 19. 4. 2010 it has rained in Paris' on 20.4. 2010 the two utterances express the same semantic content. Lasersohn has defended the same view concerning cases of faultless disagreement such as the one laid out above that involves the predicate 'fun.' According to his view, two people that disagree about whether roller coasters are fun are in disagreement about the same semantic content; they disagree about the proposition that roller coasters are fun.

2.3 How to Choose Between CT and RT

Both CT and RT are possible means of modeling context dependence and which one is favored

seems to primarily depend on the stance that one takes towards faultless disagreement about the same semantic content. (Of course, versions of CT and RT also differ with respect to the class of expressions that are considered context-dependent or assessment-relative respectively.) As it seems, the semanticist only needs to follow a simple recipe in order to determine which notion of context dependence is at play. First, he devices a context-shifting argument as laid out—and criticized—by Cappelen & Lepore (2005). One and the same sentence is considered with respect to two different utterance situations, and if it has false in one utterance situation and true in another, the sentence must contain one or more context-sensitive expressions. Once these expressions have been identified, the semanticist has to check whether it makes sense to regard the semantic content as being stable or not with respect to the variation in truth-value, i.e. whether there can be genuine cases of faultless disagreement about utterances containing the expression or not. If the answer is Yes the missing ingredient is put into the CE and if the answer is No the missing ingredient is put into the context parameter.

3. Shortcomings of the Context-as-a-Parameter View

As I will lay out now, both CT and RT lead to a number of troublesome descriptive inadequacies when they are formally modeled by regarding contexts and CEs as parameters that contribute a semantic value to the truth-conditions of the whole utterance.

3.1 The First Shortcoming: Context does not determine semantic values.

The truth-conditional content of most if not all indexicals, including the ones that Perry (1997) has dubbed ‘automatic’ indexicals, is not solely determined by ‘external’, objective, and measurable features of the context of their utterance (cf. Recanatí (2008, 173) and recently Mount (2008)). For example, a use of ‘now’ does not denote the time of utterance, but rather a time-interval of which

the time of utterance is a subinterval. When someone says 'Now is lunchtime', this does not mean that the time of utterance is the same as the entire lunchtime interval. The indexical in this sentence only expresses the constraint that lunchtime is starting from the time of utterance and includes the time of utterance. As already laid out by Fillmore (1997) in his *Lectures on Deixis*, indexicals can have broad or narrow readings to a varying degree.[2] To give an example, in an utterance of 'Taking into account inflation, salaries are not higher now then they used to be' the indexical 'now' can have a very broad reading and denote a time interval ranging from days, months, years, up to a decade. To give another example, a use of 'here' denotes a place in part of which the actual utterance took place, but whose boundaries are larger and can only be inferred on the basis of contextual information. 'It's cold here' uttered in a hut near Juneau, Alaska, can mean that it is cold on earth in comparison to the temperatures of the sun, that it is cold in Alaska, cold in Juneau, cold in the region near the hut, cold outside the hut, cold inside the hut, and so on. A use of 'we' does not denote a group of people *given by* the extralinguistic context, it denotes a group of people that the speaker has in mind or that can be vaguely determined by the social and linguistic conventions at charge within the broader, socio-linguistic context of the conversation. In 'We have achieved so much' the first-person plural indexical might denote all intelligent beings, all carbon-based beings, the members of a nation, a tribe, a family, a company, or of any other social institution, the speaker plus addressees, and so forth. Perhaps even the referent of 'I' is not uniquely determined by context when persistent tokens, figurative speech, or speaking on behalf of someone else is taken into account.[3] All of these indexicals semantically depend on the deictic center I-Here-Now (Bühler 1934) as a common, objective, and given 'core.' However, this core does not fully determine the semantic value of the use of the respective indexical. This even holds for indexicals of the 'basic set' (Cappelen and Lepore 2005), whose context-sensitivity is uncontroversial.

2 These lectures were originally given in 1971 at the summer linguistics program of the University of California, Santa Cruz and subsequently published in 1975 by the Indiana University Linguistics Club. Many examples of broad and narrow readings of indexicals can also be found in Rast (2007, 142-154).

3 Many examples of such uses of 'I' can be found in Smith (1989), Nunberg (1993), and Mount (2008).

A broader class of expressions may be called ‘contextuals.’ Expressions such as ‘fun’ or ‘tall’ extend the basic set and have either been considered context-sensitive or assessment-relative (or both). Contextuals do not even semantically depend on an objective core, or otherwise it should be possible to come to an agreement about this objective contextual ingredient as easily and uncontroversially as people tend to come to an agreement about respective features of the deictic center when an indexical is interpreted in a canonical, face-to-face utterance situation. Consider once again typical uses of indexicals. Usually it is not a problem in conversations how to identify the speaker of an utterance, how to find out where he is located and in what direction he is facing. These factors can be determined on the basis of perceptual evidence in canonical utterance situations. Likewise, the time of utterance can easily be measured by using a watch. Contextuals are clearly not on a par with indexicals with that respect. Missing ingredients of contextuals need to be inferred from various hints given by the previous conversation and the utterance situation. Perhaps per default the speaker of the respective utterance is considered the assessor of ‘fun’, but the speaker might also have someone else in mind or might want to talk about fun for everyone in a given group or for everyone in general. Certain comparison classes might be more adequate in a given situation than others, but this is always a matter of interpretation. Likewise, consider a use of ‘tall’ in an utterance of ‘Alice is tall’ in canonical face-to-face communication. While the persons present in the given utterance situation *might* constitute a contextually salient comparison class, this need not be the most salient comparison class at a given occasion. If the conversation so far has been about the swimming team of a certain university, somebody has remarked that you shouldn’t be too tall for being a good swimmer, and everybody knows that Alice is in the swimming team, it is natural to assume that the speaker meant ‘tall for a swimmer’ and perhaps additionally intended to implicate that Alice is a bad swimmer—or that Alice is a good swimmer and it is therefore not generally true that you shouldn’t be too tall to be a good swimmer. There are more or less adequate interpretations, but in none of these examples does it makes sense to stipulate an objective,

perceivable or measurable feature that serves as a basis for interpreting the expression. While all genuine indexicals semantically depend on perceivable or measurable features of the deictic center, which can be characterized either by specifying an indexical's character in a two-dimensional logic or by providing an explicit reference rule, the linguistic meaning of contextuals does not explicitly express such a dependence. A reasonable lexicon entry for 'tall' must specify that tallness can only be determined with respect to a comparison class but no explicit recipe is given how to determine this comparison class. Perhaps a lexicon entry for 'being ready' needs to indicate that you can only be ready for something,[4] yet there is no way to determine for what someone is ready directly on the basis of evidence in a concrete utterance situation as it is for example possible to perceive the place and body alignment of the speaker of an utterance.

The context-as-a-parameter view treats context as something that fixes or determines missing semantic values. The missing value is *stored* in a context parameter and *provided* by the context. While this is at least problematic in case of genuine indexicals, since in most if not all cases only an objective 'core' of their actual truth-conditional contribution is given in that way, this view is outright inadequate when a broader class of contextuals is taken into consideration (cf. Penco (1999)).

3.2 The Second Shortcoming: A fruitful notion of context must account for more than the mere variability of extensions from assessor to assessor.

It seems that the recent revival of relativist positions has at least in part been accelerated by the contextualism vs. minimalism debate between authors like Recanatí (2004), Bach (2005, 2007), Borg (2004), and Cappelen & Lepore (2005)—just to name a few. This is somehow surprising,

4 Perhaps in this case it is more plausible to assume a minimalist, non-relational lexicon entry as a unary predicate than in case of being tall, but let's not delve into these minimalist intuitions further. If there is disagreement in a particular case, it should always be possible to come up with an example that is more clearly relational, unless one shares the strong minimalist intuitions of Cappelen and Lepore.

because the debate on one hand marks a general shift towards contextualism, but on the other hand RT based on two-dimensionalism is in a sense 'more' semantic than CT, as it relativizes truth and falsity of semantic content to an assessor. Sensitivity to an assessor can be put into the context parameter in CT or into the CEs in RT, and the difference between the two positions boils down to whether there is disagreement about differing content with the same underlying linguistic meaning or about the same content with the same underlying meaning. However, when somebody interprets an utterance containing indexicals or contextuals, he *interprets* these expressions just as any other expressions on the basis of his current assumptions and their linguistic meaning. In other words, genuine interpretation always operates on the level of linguistic meaning and not on the level of semantic content derived in dependence of factors external to the interpreter's web of beliefs. So assessment-relativity is not a substitute for interpretation.

By stipulating that the truth or falsity of the semantic content of a sentence like 'roller coasters are fun' is relative to an assessor, RT elevates the assessor himself to the status of a partial, objective truth-maker, whereas in fact he is, as the name implies, the person who assesses the truth or falsity of a sentence in a context. This is a conceptual flaw of RT only, whereas the previous point is a conceptual flaw of both parameter-based CT and RT.

3.3 The Third Shortcoming: From a technical perspective, parameterized contexts and assessment-relative CEs are of limited use for modeling interpretation.

Here is a strategy for maintaining parameter-based CT or RT in light of the previous criticisms: Introduce an interpreter or assessor into the context or CE respectively for each interpretation of the sentence by some agent. This is in fact what RT does with a truth-relativist twist. Unfortunately, such an account is not very helpful for describing or explaining the Hows and Whys of interpretation. Single context parameters and CEs prima facie do not represent an agent's

assumptions and beliefs according to which he interprets linguistic meaning in a certain way in a given situation. It is impractical, albeit not impossible, to squeeze beliefs and assumptions into parameters. Lewis (1980) enriched indices of modal logics in this way, yet as a general remedy to the before mentioned problems such an approach remains unsatisfactory, because a suitable representation of assumptions and beliefs must be able to represent ignorance. In order to represent the possibility that an agent considers a certain proposition compatible with his assumptions or beliefs respectively without actually assuming or believing it, a *set* of reference points (states, possible worlds, situations) or a more fine-structured entity is needed. Moreover, it would be generally desirable to be able to express interesting features of interpretation when the latter is understood as an ideally-rational reasoning process. For example, being able to describe the case when an agent draws inferences from a set of assumptions about the utterance situation and about what has been said earlier seems crucial for explaining how an ideally-rational agent arrives at a certain interpretation. In addition to this, revisions and the updating of assumptions and beliefs need to be modeled whenever the dynamics of discourse is taken into account. Finally, all of these operations should preferably be available within a logical language with multiple agents. In contrast to this, a parameter is something that is kept fixed most of the time and barely changes. Although it is feasible to enrich parameters in various ways, this strategy seems to be rather limited for the modeling of the *reasoning processes* which allow an ideally-rational agent to arrive at an interpretation on the basis of his current beliefs and assumptions and for the modeling of the subsequent checking and updating of these beliefs and assumptions by the semantic content that results from the interpretation process.

From this mainly external critique it ought not be concluded that the whole idea of two-dimensionalism should be given up, though. An agent interprets nonindexical expressions in an utterance on the basis of his assumptions and beliefs about the topic of the utterance, i.e. the situation described by the utterance or what the utterance is about, and indexicals on the basis of his

respective assumptions and beliefs about the utterance situation. To see this, consider an utterance of 'We have believed for a long time that steel is strong enough.' The indexical 'we' semantically depends on the speaker of the utterance and so the utterance situation has to be taken into account. However, a reasonable interpretation of the contextual 'enough' does not need to be based on assumptions or beliefs *about the utterance situation*, given that contextuals do not generally depend on the deictic center as I have laid out above. Topic situation and utterance situation only coincide when the present tense is used and no other modalities are at play, but even this assumption may be dropped for some purpose by the same token as truth-relativists drop the initial coupling between features of the context and the CE. In any case, for what steel is strong enough must be inferred by a particular agent on the basis of what he believes that the utterance is about—on the basis of his believes about the topic situation. Is the utterance about building a construction to support the roof, an airplane, a vault, etc.?

Thus, two situations or corresponding more complex reference points are required, although only subjective opinions about them are relevant for the interpretation process itself. While the interpretation model of context is designed for taking this into account, the version I will lay out in the next section is relatively modest. With a number of small technical changes it is possible to implement IP within a two-dimensional logic. However, conceptually the resulting account differs considerably from the versions of CT and RT outlined so far.

4. The Interpretation Model

4.1 Context as Assumptions: Context is a set or structured ensemble of assumptions and beliefs.

The above problems have been known for a long time in AI research community, where starting with McCarthy (1993) a number of powerful logics of context have been developed to circumvent them. See for example Buvac (1995), Giunchiglia (2001), and Serafini (2004). Logics with reified

contexts are viable alternatives to modal logics with parameters, but they don't seem to be mandatory for the modeling of interpretation. According to the more conservative approach that I will sketch in the following paragraphs, the interpretation model can also be implemented on the basis of a two-dimensional modal logic as long as certain conceptual and technical changes are made. What follows is a rough sketch of this approach.

First of all, as in von Stechow/Zimmerman (2005) contexts and CEs need to have the same structure, since they now represent situations. One parameter represents the utterance situation and the other one represents the topic situation, which differs from the former when it has been shifted by a traditional modal operator for the alethic modalities, ordinary rational belief, and tenses. Each of these situations has its own domain that can be retrieved by a domain function. Moreover, situation parameters ought to be small in the sense it doesn't follow from the fact that *p* is not the case at a situation *s* that *not p* is the case at *s*. This can be achieved by using a partial logic, but I will not go into these details here.

Given that parameters represent situations, an assumptive interpretation operator may be defined for each agent, which shifts both the context parameter and the CEs to those that are compatible with what the respective agent assumes in a given interpretation situation. Hence, not only nonindexical but also indexical expressions in the scope of this operator are evaluated with respect to the agent's current assumptions. For example, in this view David's assumptive interpretation of 'It's cold here' will turn out true in a model if and only if it is cold at the place that is compatible with what he currently assumes *about* the utterance situation at a time interval of which the time of what he assumes to be the utterance situation is a subinterval. This condition is only fulfilled when 'here' denotes the same place in all situations compatible with what David assumes about the utterance situation, and in this particular case utterance and topic situation coincide.[5]

To give another example, consider an utterance of 'Alice is tall.' Let the predicate *Tall(a)* be true

5 Perhaps this condition should be relaxed to take into account sufficient similarity between places. There are many interesting details that cannot be further addressed here for lack of space.

with respect to an utterance situation *u* and topic situation *t* if and only if for most *x* in the domain of *t*, *height(x)<height(T(a, u, t))*, where *T* stands for term interpretation with respect to an utterance and a topic situation. This lexicon entry for 'tall' is admittedly naïve, but suffices for the present purpose. Let $R_i(x, y)$ be a binary accessibility relation over situations that represents compatibility with David's interpretative assumptions in a given interpretation situation *i*. Ignoring the temporal dimension for simplicity, his interpretation of *Tall(a)* is then true if and only if for all *t', u'* s.t. $R_i(t, t')$ and $R_i(u, u')$ most *x* at the domain of t' are such that *height(x)<height(T(a, u', t'))* provided *T(a, u', t')* denotes the same individual in *t'* (*u'* being redundant in this case), false or, in a partial logic, undefined otherwise.[6] Whether or not Alice is tall according to David's interpretation does not depend on objective, measurable features of the utterance or topic situation, but only on what he assumes about them.

While it is trivial to device a model in which 'Alice is tall' comes out true with respect to a particular actual topic situation, models may also be restricted to disallow such cases. Nothing prevents us from stipulating that *Tall(*x*)* is false or undefined with respect to any *actual* topic situation. This is a key feature of the interpretation model of context. Someone's assumptive interpretation of an utterance can be meaningful, even though the respective sentence taken for itself is false or has no definite truth-value with respect to the actual utterance and topic situations, as these may be underspecified with respect to the given indexical or contextual.

4.2 Full Interpretation: Even when Gricean processes are not taken into account, assumptive interpretation as an intensional operator is only a small piece in the puzzle of how to model interpretation.

It is important to stress that the notion of assumptive interpretation does not represent interpretation

6 Notice that *a* must be a nonrigid constant in order for this to work. In the above case, *a* is not indexical and thus *T(a, u, t)=T(a, u', t)* for any *u, u'*. The dependence on the utterance situation is redundant in this example, since for simplicity tense has been ignored.

in the broad sense and only provides some basic means to start reasoning about full-fledged interpretation. An intensional interpretation operator works in a straightforward way whenever a missing ingredient can be derived naturally from the underlying position about what parameters are from a metaphysical point of view. Situations have an associated domain, time, place, and world. Assumptive interpretations of quantifier domain restrictions and other domain-sensitive expressions like 'tall', temporal indexicals, and spatial predicates and indexicals can therefore be modeled in a natural way by such an operator. The interpretation of many other expressions, contextuals in particular, depends on much more complicated, heuristic processes, though. For example, before someone is ready to leave a party, he will usually say goodbye to many people, grab his jacket, make sure that he hasn't forgotten anything, and so forth. So when someone utters 'John is ready', an interpreter might infer that he is ready to leave the party and drive home in the basis of a vast number of assumptions and beliefs about the topic situation and a whole background theory about the usual behavior of people in party situations. It would be inadequate to ignore this inferential process and somehow encode the 'ready for what?' part in contexts or CEs. In light of this example, the suspicions of semantic minimalists like Cappelen/Lepore (2005) and Borg (2004) against the idea of continually enriching parameters with all kinds of additional values are certainly justified.

Figuring out that someone is about to leave a party is of mind-boggling complexity. Nevertheless, the interpretation model of context seems to generally fare better with this task than simple parameter-based approaches. Assume, following Bach (2005, 2007b), that an utterance of 'John is ready' expresses an incomplete propositional skeleton. This might for example be represented as an open higher-order logic formula *ReadyFor(j, X)*. While the missing *X* might not always be just one thing, at least the default interpretation of such an incomplete content seems to be the one where the missing arguments are bound by an existential quantifier (cf. Borg (2004)). After all, it is not likely that the speaker wants to talk about several ways in which John is ready by making a single utterance. A default completion of the utterance is thus $\exists X$ *ReadyFor(j, X)*. Given that, the question

is how an interpreter arrives at a more specific interpretation. If this is bound to be explained by broadly-conceived logical means, which is unlikely to reflect what's going on in an interpreter's brain, then it seems that there is a lot of default reasoning involved in finding a particular value for *X*—complex reasoning that is relative to an interpreters assumptions and therefore must somehow be explained on the basis of an assumptive interpretation operator or by similar means. Given a rich enough set of assumptions a particular value for *X* might for example be found as follows.

Background knowledge: When someone is about to do something, he is usually ready to do it (whatever that means). When somebody is at a party, says goodbye to everyone, grabs his jacket and checks that he hasn't forgotten something, he is usually about to leave the party.

Particular assumptions of the interpreter: This is a party situation, John is at the party, he says goodbye to everyone, he grabs his jacket, etc., and so (unless there is specific evidence against the default conclusion) he is about to leave the party.

Particular inference: *ReadyFor(j, Leave(j, p))* unless there is further evidence that would block the second default conclusion.

The goal of the interpretation model of context is to explain such inferences as defeasible deductions from the assumptions of a given interpreter. In many cases the path from less specific to more specific content is not deterministic. It is easy to conceive sets of assumptions that allow for alternative substitution instances for the existentially-bound variable in the above example. This is how it should be, since when interpreting an utterance one rarely arrives at just one possible way to understand it. For example, when in a previous conversation people at the party, including John, were talking about their plans for holidays, another, perhaps more complicated trail of thought

might lead to the inference that 'John is ready' should be taken as 'John is ready for going on a holiday' in the given situation. In cognition, possible interpretations are continually produced and ranked according to their plausibility and the one that has highest plausibility at a given time is chosen. This view has been defended by relevance theorists like Sperber & Wilson (1986) and Sperber (2006) for a long time using slightly different terminology. While it is intuitively appealing, it unfortunately poses considerable obstacles to a logically-oriented implementation that is supposed to describe and explicate corresponding reasoning processes of ideally-rational agents. Not only default reasoning, but also degrees of plausibility are needed according to which possible interpretations are ranked.[7] In combination with what has been said earlier, the overall technical apparatus required for applying the interpretation model of context to less than trivial examples is considerable. Perhaps this is one of the reasons why philosophers have resorted to seemingly simpler explications of dealing with the semantics and pragmatics of indexicals and contextuals—explications that are not based on an interpreter's assumptions. However, what has been said in the previous paragraphs has hopefully made it clear that oversimplifying interpretation processes causes more problems than it solves. The resolution of context dependence, even that of many automatic indexicals, is an interpretation process, and full interpretation *is* a complex phenomenon.

4.3 The Role of Speaker Intentions: Albeit playing a crucial role in the process of interpreting an utterance, speaker intentions do not generally fix or determine missing ingredients.

From a more Gricean perspective it may be argued that a speaker's intentions fix the missing ingredient of indexicals and contextuals. Let's call this the intention-based model of interpretation (IMI). I submit that this view is too speaker-centric and not adequate in general.

Before presenting an argument against IMI, a short disclaimer is necessary. When talking about

7 There are numerous ways to formally represent belief that comes to a degree and drawing inferences thereof, see Paris (1994) and Halpern (2003) for an overview. However, these do not scale well to first- and higher-order languages and there are still many philosophical and logical problems to be solved in this domain.

interpretation, so far always interpretation of literal meaning, i.e. the part of meaning of an utterance that is derivable from a shared lexicon, has been meant. There is no doubt that an interpreter *also* attempts to figure out what the speaker wanted to convey by making the utterance, but the above notion of interpretation is not devised for describing this Gricean process. Completing missing ingredients of indexicals and contextuals occurs at the locutionary level. Or, remaining terminologically more neutral, one may say that contextual completion is a process that occurs at a different level of content than what happens when an interpreter is trying to figure out what the speaker meant to convey. For example, 'John is ready' uttered by Bob might at the illocutionary level be interpreted as an indirect command to Alice. When Bob and Alice believe that she has previously asked John to drive her home, Bob might intend to convey that Alice ought to get ready to leave the party, too. When figuring out what the speaker meant to convey by an utterance, an interpreter needs to take the speaker's communicative intentions into account. However, as I have laid out in the previous sections, utterances already need to be interpreted at the locutionary level and at this level a speaker's intentions play a different and perhaps less important role.

Suppose Bob has 'John is ready to leave the party' in mind when he utters 'John is ready.' In some sense of 'to intend something' one might say that Bob intends Alice to grasp the meaning of 'John is ready to leave the party.' Suppose that Bob wants her to grasp that John is ready to leave the party in order for her to realize that he wants to give her the indirect command to get ready herself.

I submit that this content, the proposition that John is ready to leave the party, does not necessarily represent what is said by Bob's utterance of 'John is ready.' Assume, for example, that Alice has just been talking with Carol about how stressful work is at the moment and how ready they are to go on a holiday. Not having heard the previous conversation, Bob suddenly chimes in and utters: 'John is ready.' Even if Bob wants to convey at the illocutionary level that John is ready to leave the party, it would be awkward to say that 'John is ready to leave the party' is the one and only 'correct' interpretation of his utterance at the level of what is said in the given situation—as opposed to what

Bob says by making the utterance. Alice and Carol have good reasons to assume that he meant 'John is ready to go on a holiday' and no reason to assume he meant something else, and Bob has misjudged the background assumptions of the other discourse participants. Already at the locutionary level, there are two distinct interpretations with two different associated adequacy criteria at play in this case. On one hand, Bob intends to convey the proposition that John is ready to leave the party at the locutionary level in order for Alice to grasp his indirect command at the illocutionary level. On the other hand, Alice and Carol's interpretations are more adequate than the interpretation Bob has in mind in this case, because their interpretations rest on the most salient shared background assumptions. Hence, Alice and Carol concur while misunderstanding Bob. So the adequacy of an interpretation does not only depend on whether the interpreters figure out what the speaker meant to convey, but also on the correctness of the speaker's assumptions about the background assumptions of all discourse participants in the given situation.

This is, of course, also a crucial component of the full Gricean account. I do not argue against Gricean pragmatics in general here, but only against a speaker-centric 'shortcut' version of it, according to which missing ingredients are simply fixed by a speaker's intentions. While in many types of discourse involving assertions interpreters are often if not always interested in figuring out what the speaker meant, the utterance itself gives rise to a form of content at the locutionary level that is not necessarily obtained on the basis of an interpreter's assumptions about the speaker's communicative intentions. In the above example, Alice's and Carol's interpretation is adequate and Bob's intended interpretation is deviant. So even if a speaker's intentions suffice to fix or determine illocutionary content, they do not in general fix or determine locutionary, yet contextually-completed content.

To give another example, suppose Alice and Bob are looking for a wine that doesn't cost more than $20. Being unaware of this requirement, Carol points to a bottle of wine that costs more than $100 and says 'This wine is good.' Perhaps she has tasted the wine before and wants to convey that the

wine tastes good. Nevertheless, Alice and Bob may rightfully reply that the wine is not good, since based on their shared beliefs they have good reasons to prefer a different interpretation: that the wine is good, because it might taste good and costs no more than $20—which is false if the wine under consideration costs more than $100.

In order to figure what the speaker meant literally, in this case that the wine is good in the sense of tasting good (there is perhaps more to being a good wine than just tasting good, but let's put this matter aside here), and in order to figure out the illocutionary content such as 'Buy this wine, because it is good!' the interpreters may have to take into account the speaker's communicative intentions, but in order to obtain a reasonable and in the sense laid out above adequate interpretation of the utterance itself in the given situation, they do not *have to* consider the speaker's intentions.[8] They do not have to, because it is the speaker's obligation to choose any conversational shortcut such as indexicals, contextuals, or ellipses in such a way that it is supported by the Common Ground and conforms with the Gricean maxims. IMI does not take this aspect of interpretation sufficiently into account, given that what an utterance says in a given situation can be as relevant in actual communication as what the speakers intends to convey by making the utterance.

4.4 Assumptions versus Belief: In a given interpretation situation, assumptions are generated by an interpreter's previous first- and higher-order beliefs, and the result of interpreting an utterance is in turn checked against his previous beliefs.

Being primarily concerned with modeling pragmatic presuppositions and their accommodation, Stalnaker (1978) emphasized that the Common Ground is constituted by the assumptions that discourse participants mutually share at a given time, and not by what they mutually believe. As is well-known, Stalnaker modeled these assumptions as a set of possible worlds, the 'context set',

8 Although he seems to be generally inclined towards Gricean explications (implicatures notwithstanding), Kent Bach has made very similar points. See Bach (2005).

which is updated by the proposition expressed by a new utterance and by its presuppositions when all discourse participants accept it. IP is a straightforward generalization of this view and as in case of Stalnaker's original suggestion it is crucial to take an interpreter's assumptions and not his beliefs simpliciter as the basis of modeling interpretations. As I will lay out now, these assumptions are *generated* at a given time from an interpreter's beliefs.

As Stalnaker, Lewis, and many others have pointed out, people often silently accommodate presuppositions without necessarily endorsing them. For example, if someone talks about the present king of France, an interpreter might for the sake of understanding what the speaker intends to convey not protest against the presupposition that France currently has a king and simply assume that there is one. The speaker is usually aware at least of obvious presuppositions of his utterance and interpreters can therefore assume that he believes what his utterance presupposes under the assumption that the speaker is honest and sincere.

A similar, though slightly different phenomenon occurs when context-sensitive expressions and semantically incomplete content is interpreted. Consider, for example, the above mentioned conflict between the interpretation that Bob has in mind when uttering 'John is ready' and the interpretation of Alice and Carol. If they believe that Bob has not overheard their previous conversation, they might discard their own interpretation in favor of an interpretation that they believe to be more salient *for Bob* than theirs. In other words, they might interpret his utterance on the basis of what they believe Bob believes at the current time rather than on the basis of their first-order beliefs. Adjusting one's assumptions on the basis of what one believes about what someone else believes is a crucial aspect of interpretation.

Technically, this process can be regarded as the revision of an interpreter's first-order beliefs about the utterance and topic situation respectively by his beliefs about what the speaker believes about the utterance and topic situation in a given interpretation situation. The resulting sets of situations represent what the interpreter assumes (in the interpretation situation) about the utterance and topic

situation respectively.

To see how this might work, let me take a look at an utterance of 'John is ready' again. If Bob chimes in without having overheard the previous conversation between Alice and Carol, Alice might form the belief that Bob doesn't believe that the current conversation is about her and Carol's plans for their next holiday. After revising her first-order beliefs about the topic of the conversation with her beliefs about what Bob believes, the resulting assumptions about the topic of the conversation will no longer contain the proposition that the current conversation is about going on a holiday. However, all of her previous beliefs compatible with the proposition that the conversation is not about going on holidays will be maintained after this revision process. On the basis of these assumptions she might then arrive at the interpretation that John is ready to leave the party in the way indicated above.

Hooking up belief revision theory in the tradition of Alchourrón (1985) and Gärdenfors (1988) with a multi-agent sender-receiver model of communication is not a trivial task. Iterated belief modalities and the higher-order nature of many examples pose interesting technical challenges.[9] Everything comes at a price, but the rewards of explicitly modeling the ideal reasoning processes behind interpretation in a context are correspondingly high. The pitfalls and shortcomings of parameter-based contextualism and relativism that I have laid out in the beginning of this article are avoided and the interpretation model provides the cornerstone for the description and explication of more complex pragmatic phenomena such as deferred ostension or narrative discourse.

5. Summary

The view on context that I have laid out in this article is based on the following claims. Context-sensitive expressions are interpreted by an agent on the basis of his assumptions and the linguistic

9 Recent results in Dynamic Doxastic Logic such as and Dynamic Epistemic Logic can deal with many of the problems of a more technical nature such as how to model belief change in languages that allow iterated modalities. See van Benthem (2006), Leitgeb & Segerberg (2007), and Ditmarsch et. al. (2008) for an overview. However, how to implement belief revision and contraction operations in a higher-order setting remains an open problem as to the time of this writing.

meaning that they express by virtue of what is derivable from a shared lexicon, whether or not this meaning can be true or false with respect to a context and a topic situation. Following Bach in that respect and acknowledging that utterances may literally express propositional skeletons is crucial in this view, but as laid out above a weak default interpretation can always be obtained by existentially quantifying over open argument places. Utterances are interpreted with respect to a topic and an utterance situation. Interpretations of expressions sensitive to domains, times, places, and worlds can be adequately captured by an assumptive interpretation operator, whereas the interpretation of broader contextuals like ‘ready’ or ‘enough’ depends on often complex default reasoning processes operating on the basis and on top of an interpreter’s assumptions at a given time. Context is an ensemble of beliefs and assumptions that can be generated from an agent’s beliefs in a given interpretation situation by revising his first-order beliefs by his second-order beliefs about what the speaker believes.

Literature

Alchourrón, C. E.; Gärdenfors, P. & Makinson, D. (1985), 'On the logic of theory change: partial meet contraction and revision functions', *Journal of Symbolic Logic*(50), 510-530.

Bach, K. (2007), 'Minimalism for Dummies: Reply to Cappelen and Lepore', Technical Report, University of San Fransisco, Department of Philosophy, available from URL *http://userwww.sfsu.edu/~kbach/replytoC&L.pdf*.

Bach, K. (2005), Context ex Machina, *in* Zoltán Gendler Szabó, ed.,'Semantics versus Pragmatics', Oxford UP, Oxford, pp. 16-44.

Benthem, J. v. (2006), 'Dynamic Logic of Belief Revision', Technical report, University of Amsterdam, Institute for Logic, Language and Computation.

Borg, E. (2004), *Minimal Semantics*, Oxford UP, Oxford / New York.

Bühler, K. (1934), *Sprachtheorie*, Fischer, Stuttgart, Jena.

Buvac, S. (1995), 'Formalizing Context' (FS-95-02), Technical report, AAAI.

Cappelen, H. & Lepore, E. (2005), *Insensitive Semantics*, Blackwell, Oxford.

Carcía-Carpintero, M. & Kölbel, M., eds. (2008), 'Relative Truth', Oxford UP.

Ditmarsch, H. v.; Hoek, W. v. d. & Kooi, B. (2008), *Dynamic Epistemic Logic*, Kluwer.

Fillmore, C. J. (1997), *Lectures on Deixis*, CSLI Publ., Stanford.

Gärdenfors, P. (1988), *Knowledge in Flux*, MIT Press.

Ghidini, C. & Giunchiglia, F. (2001), 'Local Models Semantics, or Contextual Reasoning = Locality+Compatibility', *Artificial Intelligence* **127** (2), 221-259.

Halpern, J. Y. (2003), *Reasoning About Uncertainty*, MIT Press.

Kaplan, D. (1989), Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals., *in* Joseph Almog; John Perry & Howard Wettstein, ed., 'Themes from Kaplan', Oxford University Press, Oxford, New York, pp. 481-564.

Korta, K. & Perry, J. (2007), Radical Minimalism, Moderate Contextualism, *in* Gerhard Preyer & Georg Peter, ed.,'Context-Sensitivity and Semantic Minimalism', Oxford UP, Oxford, pp. 94-111.

Leitgeb, H. & Segerberg, K. (2007), 'Dynamic doxastic logic: why, how, and where to?', *Synthese* **155**(2), 167-190.

Lewis, D. K. (1980), Index, Context and Content, *in* S. Kanger & S. Öhman, ed.,'Philosophy and Grammar', D. Reidel Publishing Co., Dordrecht, Boston, pp. 79-100.

MacFarlane, J. (2007a), 'Relativism and Disagreement', *Philosophical Studies* **132** (1), 17-31.

MacFarlane, J. (2007b), 'Nonindexical Contextualism', *Synthese* **166** (2), 231-50.

MacFarlane, J. (2007c), 'Relativism and Disagreement', *Philosophical Studies* **132** (1), 17-31.

MacFarlane, J. (2005), Making Sense of Relative Truth, *in* 'Proceedings of the Aristotelian Society', pp. 321-39.

McCarthy, J. (1993), Notes on Formalizing Context, *in* Ruzena Bajcsy, ed., 'Proceedings of IJCAI', Morgan Kaufmann, Chambery, France, pp. 555-562.

Mount, A. (2008), 'The Impurity of "Pure" Indexicals', *Philosophical Studies* **138**, 193-209.

Nunberg, G. (1993), 'Indexicality and Deixis', *Linguistics and Philosophy 16:1 (1993)*(16:1), 1-43.

Paris, J. (1994), *The Uncertain Reasoner's Companion: A Mathematical Perspective*, Cambridge University Press, Cambridge.

Penco, C. (1999), Objective and Cognitive Context, *in* M. Benerecetti & F. Castellani, ed., 'Modeling and using context : Lecture Notes in Artificial Intelligence 1688', Springer, pp. 270-283.

Rast, E. H. (2007), *Reference and Indexicality*, Logos, Berlin.

Recanatí, F. (2008), Moderate Relativism, *in* Manuel García-Carpintero & Max Kölbel, ed.,

'Relative Truth', Oxford UP, Oxford / NewYork, pp. 41-62.

Recanatí, F. (2004), *Literal Meaning*, Cambridge University Press, Cambridge, UK.

Serafini, L. & Bouquet, P. (2004), 'Comparing Formal Theories of Context in AI', Artificial

Intelligence Journal (155), 41-67.

Smith, Q. (1989), 'The Multiple Uses of Indexicals', *Synthese* **78**, 167-191.

Sperber, D. & Wilson, D. (2006), Relevance Theory, *in* Laurence R. Horn & Gregory Ward, ed.,

'Handbook of Pragmatics', Blackwell, Oxford.

Sperber, D. & Wilson, D. (1986), *Relevance: Communication and Cognition.*, Blackwell, Oxford.

Stalnaker, R. (1978), Assertion, *in* P. Cole, ed., 'Pragmatics', Academic Press, New York, pp. 315-332.

Stanley, J. (2002), Nominal Restriction, *in* Gerhard Preyer & Georg Peter, ed., 'Logical Form and Language', Oxford UP, Oxford, pp. 365-388.

Stanley, J. & Szabó, Z. G. (2000), 'On Quantifier Domain Restriction', *Mind & Language* **15** (2 & 3), 219-261.

Stojanovic, I. (2008), 'Talking about Taste: Disagreement, Implicit Arguments, and Relative Truth', Technical Report, Institut Jean Nicod, Paris, URL *http://hal.archives-ouvertes.fr/docs/00/28/11/02/PDF/LING146R2-talking-about-taste-1-may-2008.pdf*.

Stojanovic, I. & Predelli, S. (2008), Semantic Relativism and the Logic of Indexicals, *in* Manuel Carcía-Carpintero & Max Kölbel, ed., 'Relative Truth', Oxford UP, Oxford / New York, pp. 63-80.

Travis, C. (2008), *Occasion-Sensitivity*, Oxford UP, Oxford / New York.

von Stechow, A. & Zimmermann, T. E. (2005), A Problem for a Compositional Treatment of De Re Attitudes, *in* Gregory N. Carlson & Francis Jeffrey Pelletier, ed., 'Reference and Quantification: the Partee Effect', CSLI Publications, pp. 207-228.