

NOTICE: This is the author's version of a work that was accepted for publication in <International Review of Pragmatics>. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms, may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in INTERNATIONAL REVIEW OF PRAGMATICS, 6(1), 2014, DOI:<http://booksandjournals.brillonline.com/content/journals/10.1163/18773109-00601003>

# The continuum problem: Modified Occam's Razor and conventionalisation of meaning

Abstract: According to Grice's "Modified Occam's Razor", in case of uncertainty between the implicature account and the polysemy account of word uses it is parsimonious to opt for the former. However, it is widely agreed that uses can be *partially* conventionalised by repetition. This fact, I argue, raises a serious problem for MOR as a methodological principle, but also for the substantial notion of implicature in lexical pragmatics. In order to overcome these problems, I propose to reinterpret implicatures in terms of implicature-like effects delivered by non-inferential processes.

Keywords: lexical meaning; semantics/pragmatics; implicatures; associative processes

## 1. Introduction

In Paul Grice's framework for pragmatics, the principle called Modified Occam's Razor (MOR) is an important means to cope with the demarcation between semantic and pragmatic phenomena. Grice has employed it in order to argue for a pragmatic treatment of standard uses of certain linguistic forms: indefinite descriptions, the counterparts of logical connectives in ordinary language, etc. The assumption

is that, in the explanation of such cases, “Senses are not to be multiplied beyond necessity” (Grice, 1978: 118-119); in other words, one should not make an appeal to a special sense for those expressions insofar as a general explanation can be supplied, based on the reasonable hypothesis that the use at issue can be derived from a basic meaning through the application of some pragmatic principle.

Recently, the reliability of MOR has been called into question with arguments that converge towards a common point, that is, the difficulty to draw a straight line between fixed lexical meanings and pragmatic inferences (Bontly, 2005; Phillips, 2012; Mazzone, 2013 a): let us call it *the continuum problem*. Whatever the impact of these arguments on MOR, it could be argued that research in pragmatics is not strongly affected by them. After all, MOR is just the application of a methodological principle of parsimony to a specific domain, and everyone knows that methodological principles are appealed to only as a last resort when no direct evidence is available.<sup>1</sup> As a matter of fact, scholars in pragmatics usually appeal to a variety of more specific criteria and tests in order to assess whether the use of a linguistic form is better explained as a pragmatic inference rather than as a lexicalised meaning.

However, there are reasons to think that the continuum problem has more serious and straightforward consequences on pragmatic theories than simply making apparent the weaknesses of MOR. In fact, not only is the boundary between lexicon and pragmatics difficult to draw on methodological grounds, but there also seems to be, as a matter of fact, a continuum of phenomena where Gricean pragmatics attempts to draw that boundary. Therefore, I will claim, the continuum problem does not only affect the reliability of MOR as a methodological principle, it also imposes empirical constraints on actual pragmatic theories – especially, on theories concerning the relatively recent research field of lexical pragmatics.

The existence of intermediate cases between pure lexical coding and full-fledged pragmatic inferences is recognised by Grice himself as well as by the research tradition stemming from his work. As a matter of fact, MOR has been mostly introduced in order to manage a class of cases that cannot be easily equated to prototypical (i.e., particularised conversational) implicatures, since (and to the extent that) those cases involve standard uses of words regardless of specific contexts of utterance. On the other hand, the intuition is that those uses are not part of the conventional meaning of words either. In order to account for such cases, Grice introduces the notion of generalised conversational implicatures (GCIs), that is, genuine implicatures which nonetheless are normally conveyed by specific forms of words unless

---

<sup>1</sup> MOR is an instantiation of the more general Occam's razor, which is used in science as an epistemological heuristic, that is, a provisional way to select between competing hypotheses in case decisive evidence is lacking.

the context suggests otherwise.

To be more precise, Grice's account is finer-grained than this. He also envisages another kind of implicature, conventional implicature, which would be a case in which part of the conventional meaning is nonetheless implicated rather than said. This notion of conventional implicature is not much credited in the current debate and I will leave it aside. However, in the last decades many scholars have accounted for the existence of intermediate cases between purely conventional and purely inferential (implicated) meaning in terms of some sort of conventionalisation. This is in fact suggested by Searle (1975), who speaks in this sense of “conventions of usage” that must be kept distinct from meaning conventions – and therefore from lexical meanings. In the same vein, Morgan (1978) has introduced the notion of short-circuited implicatures, that is, implicatures that are somewhat conventionalised as a result of past use (see also Horn and Bayer, 1984). Bach (1995, 1998) has proposed that the term “conventionalisation” should be reserved for cases in which precedent serves to make a given use possible, while one should speak of “*standardisation*” – and “standardised implicatures” – with regard to cases in which precedent merely facilitates the inferential process licensing a certain use (Bach, 1995: 677). As may be seen, in all these approaches the intuition of intermediate cases between particularised implicatures and conventional meaning is accounted for in terms that are very close (as is the case for Bach's “*standardisation*”), if not identical, to the notion of conventionalisation.

A very natural cognitive interpretation of this fact is based on the assumption that senses are conventionalised (i.e., coded in individual minds as conventional meanings) as a function of regularity of use: *ceteris paribus*<sup>2</sup>, the more a word form is used with a certain sense, the stronger is the association between the word form and the sense, and the more we perceive and treat the sense as a conventional meaning of that word form. As a consequence, a word meaning can be *partially* conventionalised, and this could explain the existence of intermediate cases between particularised implicatures and fully conventional meaning. The assumption at issue has strong grounds in what we know about associative learning, and it could also have a role in the explanation of how fully conventional meaning can come to be fixed at the level of individual cognition.

It is far from clear, however, that the pragmatic accounts we considered above are compatible with such a reasonable assumption, insofar as they imply a strong discontinuity between the recovery of lexical meanings and the mechanisms delivering pragmatic inferences. In fact, while it seems plausible that word uses can be partially conventionalised, it is less than clear that the notion of implicature admits

---

2 The importance of this clause will be clarified at the end of the fourth section.

of grades, and therefore that it might complement conventionalisation as it is expected to do. Intuitively, if a certain word use is in the first place licensed by a pragmatic inference and if this occurs again and again, then that use should become progressively conventionalised. But no one, to my knowledge, has ever explained what would happen to inferential mechanisms when the implicated use becomes more and more conventionalised. Do they decline to the very same extent? This does not seem quite plausible. And postulating categories of implicatures which are intermediate between fully conventional coding and pure contextual implicatures does not solve the problem. First, as we will see, such notions are debatable: they aim to satisfy a theoretical need but with weak justifications on cognitive grounds. And second, the problem is just postponed, since one would ask again how the notion of standardised implicatures or the like might take into account the fact that uses can be more and more conventionalised. Must we presume that the need for pragmatic (though standardised) inferences declines proportionally in a graded fashion? Or else, is it the case that the more standardised implicatures become usual, the less inferential – so to speak – they are?

One point should be emphasized. I do not argue against the possibility to draw conceptual distinctions such as the one involved in the notion of standardised implicature or the like. It is quite possible that cases which are intermediate between purely conventional and purely inferential meaning can be usefully categorized together for theoretical or descriptive purposes. My argument exclusively concerns the level of the cognitive mechanisms underlying pragmatic phenomena. Specifically, what I do argue is that, if one concedes that the recovery of coded meaning is sensitive to the degree of associative strength, then conceiving of implicatures in terms of inferential processes may cause inconsistency – especially in lexical pragmatics – since it is far from clear that the alleged inferential processes can complement the associative recovery as they should.

There is, however, a natural way to account for implicature-like effects without implying any strong discontinuity between lexical retrieval and alleged implicatures. The position adopted by Recanati in lexical pragmatics is the closest to adopting this view, insofar as the very same associative mechanism is held to be responsible for both lexical retrieval and (primary) pragmatic inferencing. More precisely, in Recanati's view primary pragmatic processes are associative instead of genuinely inferential, but they have nonetheless inferential-like effects. I propose to expand this approach beyond its original boundaries: implicatures in general are reinterpreted as implicature-like effects delivered by associative processes, and this may account for their gradual decline as the use of words becomes more and more

conventionalised.

In practice, I will first introduce MOR (section 2). Then I will examine the arguments that call MOR into question and the way these arguments affect theorizing on GCIs and lexical pragmatics (section 3). In section 4, I will propose an alternative account based on Recanati's non-inferential view of lexical pragmatics. Finally, given relevance theorists' particular view of pragmatic inferences, I will separately consider their approach to lexical pragmatics in the light of the continuum problem (section 5).

## 2. The Role of MOR in Grice's Theory

Grice's central intuition is that sentences may be used to convey meanings which are different from their conventional ones, and that this depends on a human inferential ability based on rational principles, thanks to which the addressee may recover the further meaning *implicated* by the speaker when, in a given context, she utters a sentence endowed with a given conventional meaning. With a famous example (Grice, 1989: 32), let us consider the following conversational exchange between two interlocutors A and B, with A “standing by an obviously immobilized car”:

(1) A: I am out of petrol.

B: There is a garage round the corner.

B's answer implicates that, as far as she knows, A will find the garage open and it will have petrol. But, obviously, these implicatures are not part of the conventional meaning of B's sentence: they are rather suggested by a consideration of the presumable communicative goal of B's utterance given both its conversational and physical context.

Although prototypical pragmatic inferences concern – just as in the above example – the further conclusions to be drawn from complete sentential meanings (in the light of the context), Grice has also taken into account a number of cases where pragmatic processes are required in order to determine the meaning that single words (or phrases) contribute to complete sentence and utterance meaning. On the one hand, these cases include disambiguation of ambiguous words and indexical reference assignment. However, as Carston (2010) has correctly emphasised, it seems as if Grice thought of disambiguation and

reference assignment in terms of mechanical processes of contextual best fit, rather than genuine processes of rational inferencing based on a consideration of the speaker's communicative intention. But there are also cases in which, according to Grice, the contribution of single words or phrases to utterance meaning is determined by means of implicatures, that is, genuine pragmatic inferences. Examples of this are indefinite descriptions and the counterparts of logical connectives in ordinary language. For instance, Grice (1989) invites us to consider an utterance of the following sentence:

(2) A man came to my office yesterday afternoon.

As a general case, we would be surprised to discover that the man was the speaker's husband since the indefinite phrase "a man" suggests that the speaker is not intimately related to that person. Grice's proposal is that this occurs as the result of a specific kind of pragmatic inference, that he calls *generalised* conversational implicature (GCIs). In contrast with *particularised* conversational implicatures, where the implicated meaning is made to depend on the specific context of utterance, GCIs are such that "the use of a certain form of words [...] would normally (in the absence of special circumstances) carry such-and-such an implicature or type of implicature" (Grice, 1989: 37). By describing the meaning contributed by the phrase "a man" as the result of an implicature, Grice intends to exclude that it is part of the fixed lexical meaning of the expression; in fact, the implicated meaning may be cancelled in case the implication is in contrast with co-textual or contextual information.

As to logical connectives and their counterparts in ordinary language, let us consider the following example:

(3) I took off my shoes and got into bed.

In this sentential context, the use of "and" appears to imply a temporal priority which is not part of its truth-functional meaning, as is shown by the fact that (3) and (4) have the same truth-functional meaning:

(4) I got into bed and took off my shoes.

Again, Grice's proposal is that temporal priority is pragmatically implicated, so that we do not need to

assume either that “and” in the ordinary language has a different meaning than its logical counterpart, or that it is ambiguous between two (or more) different meanings.

Another well studied case in point is that of scalar implicatures, that is, pragmatic effects due to “implicational scales”. These are sets of lexical items that may be ordered in accordance with their informative value, so that the use of a less informative item (a weaker value) on the scale implies that stronger values on the scale do not apply (Horn, 1972, 1984; Gazdar, 1979). For instance, in the sentence:

(5) John ate some of the cake

the presence of “some” suggests that John did not eat all the cake, since otherwise the speaker should have chosen the more informative “all”.

In sum, in all the cases exemplified in (2)-(5), the claim is that the linguistic forms at issue convey meanings which need not be, and in fact are not, lexicalised. They would rather be obtained as the result of pragmatic inferences, though of a special kind: implicatures which do not depend on specific contexts of utterance, but are normally conveyed by certain forms of words. It is mostly with regard to this sort of cases that MOR displays its function. As far as implications of the standard sort – exemplified in (1) – are concerned, there is a clear intuition that the implicated meaning is different from the conventional one and in fact produced by means of a pragmatic inference.<sup>3</sup> On the contrary, in the examples (2)-(5) there is no such clear intuition. On the one hand, the allegedly implicated meanings seem to share with lexical meanings important features – independence from contexts of utterance; standard association with specific forms of words. On the other hand, they are nonetheless claimed to be *implicated* meanings in that they share important features with prototypical (i.e., particularised conversational) implicatures: in particular, they appear to be both calculable and cancellable (I will turn to calculability below). It is precisely because the cases in point cannot confidently and undoubtedly be classified either as lexical meanings or as implicatures, since they exhibit features of both, that Grice feels the need to propose a methodological principle such as MOR. In the absence of strong evidence supporting one hypothesis over the other, Grice suggests, it could be advisable to rely on a simple principle of parsimony: we should refrain from explaining some use of an expression by appealing to specific conventions insofar as a more general explanation – based on general principles of rational

---

3 Recanati (e.g., 2004: 44), with his “availability principle” and “availability condition”, has importantly contributed to emphasize the role of this intuition.

communication – is available. In the light of this principle, then, the cases at issue would be better thought of as a special kind of implicature.

Now that we have presented the theoretical background which motivates MOR, let us consider some recent objections that have been made to this principle.

### 3. Arguments Against MOR

Bontly (2005) has summarised three main traditional arguments against MOR, and he has himself pointed to a further difficulty which has received little attention to date. Although his real aim is to show that those objections can be overcome, I intend to suggest that his attempt to defend MOR is flawed.

The first objection considered by Bontly is that there is at least a class of linguistic phenomena in the light of which the commitment of MOR to parsimony appears misleading, since it would lead to wrong predictions. The linguistic phenomena at issue are dead metaphors.<sup>4</sup> It is generally held that one-time metaphorical uses can be conventionalised as a consequence of repetition, so that a word form may come to have a secondary meaning “that is pragmatically predictable and yet fully conventional” (Bontly, 2005: 296). Pragmatic predictability is held to be, under the name of calculability, one of the characterising features of implicatures. What MOR prescribes is that whenever a word use is calculable it is more parsimonious to conclude that it is in fact calculated rather than lexically recovered. However, dead metaphors are clear counterexamples to this conclusion, since everyone agrees that they have conventional meanings which can be lexically recovered, in spite of the fact that those meanings are in principle pragmatically calculable. Therefore, parsimonious as it can be, the inference from being calculable to being calculated does not hold in general: word uses can be calculable, and nonetheless conventionalised.

While this first argument concedes that it is parsimonious to infer that *calculable* meanings are also *calculated*, though raising the objection that this inference does not hold in general, the second argument calls into question the very notion of parsimony at issue: “implicature accounts look, at least superficially, to multiply inferential labor [...]. Thus there are tradeoffs involved, and the account which

---

4 Bontly observes that versions of the argument from dead metaphors have been proposed by Walker (1975), Morgan (1978), Sadock (1978), Reimer (1998) and Devitt (2004). See also Devitt (2007).



is semantically more parsimonious may be less parsimonious all things considered” (Bontly, 2005: 297). What is at issue here is cognitive, not epistemological parsimony. Specifically, the point is that, just as cognitive linguists have repeatedly pointed out, cognitive parsimony on the storage side is inversely correlated with that on the processing side (Barsalou, 1992: 180-1; Croft and Cruse, 2004: 278; Goldberg, 1995: 74). This consideration has been taken as supportive of the view that information can be redundantly stored at different levels of linguistic taxonomies, so that it can be directly recovered from linguistic forms at any level rather than inferentially derived from superordinate levels. Although at first sight redundant storage might not appear a rational option for cognitive systems, the possibility to recover information non-inferentially might be instead a valuable gain on the side of processing effort. In the same line, by coding multiple meanings rather than employing pragmatic implicatures one could get parsimony on the processing side.

To be more precise, Bontly (2005: 297) observes that postulating an ambiguity might not dispense with inferences entirely.<sup>5</sup> Insofar as there are multiple meanings, the addressee needs to take into consideration contextual clues in order to determine which of those meanings is actually intended in a circumstance. Now, if selection of contextually relevant meanings requires inferential processing, then conventionalising a specific use of a word may lower the complexity of inferential processing, but not remove it entirely. Be that as it may, even in this case parsimony in storage and in processing are inversely correlated, and then the argument from parsimony *of storage* is far from being conclusive.

The third argument bears upon the epistemological status of MOR proper, that is, its status as a methodological device to select between competing hypotheses. The point is that methodological virtues in themselves have little impact on genuine ontological issues. On an ontological ground, the burden of proof is probably on the enemies of ambiguity: “Given that ambiguity is hardly rare, then, one wonders whether a semantic theory ought really to minimize it [...]: perhaps we ought to assume (*ceteris paribus*) that *every* regular use of an expression represents a special sense” (Bontly, 2005: 299-300, his emphasis).

In his attempt to show that MOR can be defended from the received criticisms, Bontly starts from this last epistemological point. He makes an appeal to Sober's (1988) claim that arguments from parsimony may be cogent after all, provided that they are employed not as domain-general principles of

---

5 Here and elsewhere, Bontly's use of the term “ambiguity” is potentially misleading. A word is said to be ambiguous when it has a plurality of possibly unrelated meanings, while the debate on MOR concerns cases in which meanings can be rationally inferred from one another and are, therefore, related to each other. These are better referred to as cases of “polysemy”.

scientific reasoning, but instead with reference to domain-specific assumptions about the processes that generate the phenomena under study. In practice, Bontly's argument is that there is a tendency to parsimony that is specific to semantic acquisition. In fact, it is a well established psychological observation that children have difficulties learning homonyms,<sup>6</sup> and this suggests that “language acquisition is semantically conservative: *children will posit new meanings for familiar words only when necessary*” (Bontly, 2005: 302, emphasis mine). This conclusion, in Bontly's opinion, is precisely the sort of domain-specific assumption that makes MOR a reasonable principle for theory-choice in semantics: since subjects resist positing new meanings (except “when necessary”), then there is a specific cognitive justification for semantic parsimony.

Unfortunately, the argument crucially depends on how one interprets the words “when necessary”.<sup>7</sup> Bontly apparently assumes that the right interpretation is: “when there is no pragmatic inference by which an addressee could calculate the intended meaning” or, more synthetically, “when meanings are not calculable”. Only by this interpretation would the argument from acquisition speak in favour of MOR: if subjects resist positing new meanings unless these cannot be calculated, then we should not attribute stored meanings to subjects when those meanings are calculable. But there are other possible interpretations of “when necessary”, for instance, “when forced by repetition”. In this case, observations from acquisition would speak in favour of a wholly different conclusion, that is, the conclusion that we should better not attribute stored meanings unless when those meanings are repeatedly used (in association with a phonological form). Since Bontly simply assumes the former interpretation without providing any reason for it, his argument is a *petitio principii*: that new meanings are resisted (and eventually not stored) *in case they are calculable* is what should be shown, while it is in fact tacitly

---

6 The same terminological warning of note 5 applies here to the term “homonym”: this specifically refers to multiple meanings which are unrelated to each other. Bontly apparently thinks that the cited evidence licenses conclusions concerning polysemy as well. As noted by an anonymous referee, this inference is far from warranted: children could resist homonymy and yet “be perfectly happy with words being polysemous”. Although this is in itself a serious problem for Bontly's argument, I propose a more general criticism which applies also in case it would turn out that children resist polysemy as well.

7 The evidence cited by Bontly, for instance Mazzocco (1997), only shows that young children (3- to 5-year-olds) tend not to use already known words for new referents, while they can easily learn new words for those referents. In the light of those studies no conclusion can be drawn as to the reasons why children are so “parsimonious” in the former case, or the conditions at which such parsimonious tendency would be overridden. In particular, there is nothing suggesting that children are always and only parsimonious when meaning can be calculated.

assumed by the argument.

Not only would Bontly need an argument in order to adopt his interpretation of “when necessary” without begging the question, but he also needs a way to manage the case of dead metaphors. As we saw, dead metaphors seem to show that meanings can be stored even when they are calculable, and this seems to be a strong case against the thesis that children in particular and people in general do not store new meanings except when these are not calculable. In fact, Bontly attempts to provide an argument for explaining away the case of dead metaphors. On the basis of his developmental account of parsimony, Bontly proposes the following analysis of what he assumes to be an instance of dead metaphor, that is, the metaphorical use of “incense”:

It may be, therefore, that children at the relevant developmental stage [...] lack the understanding necessary to predict that “incense” could be used to mean *to make or become angry*. They might not realize, for instance, that “incense” means *an aromatic substance that burns with a pleasant odor*, and even those who do probably lack the general background knowledge necessary to appreciate the metaphorical connections between burning and emotion. Either way, the metaphor would be dead to children, forcing them to learn that use the same way they learn any arbitrary convention. (Bontly, 2005: 309)

This seems to be a reasonable analysis of how the word “incense” is acquired by children. Bontly apparently thinks that a general lesson can be drawn from this example, and it is that children are forced to store metaphorical meanings only insofar as they cannot calculate metaphors for lack of the knowledge required. However, one wonders whether his analysis generalises to any case of presumed dead metaphors: does it apply, for instance, to “legs of tables”, or to “hiding behind someone else” (in its figurative meaning), and so on and so forth? Since there are a number of cases in which children appear to have the background knowledge necessary to appreciate the relevant metaphorical connections, then at least in those cases there is still an argument against MOR and for storage of calculable meanings, unless one gives independent reasons for believing that metaphorical meanings are stored in “incense” cases but not in “leg of the table” ones. Since, on the contrary, Bontly gives no such reason, he is again begging the question.

To sum up, there is circularity both in Bontly's developmental argument for MOR, and in his specific appeal to the same argument in order to manage the case of dead metaphors. But Bontly has also made an interesting observation that helps to shed light on where the real problem with MOR lies.

As we saw, there is a third traditional argument against MOR (besides the argument from dead metaphors, and the argument from the distinction between epistemology and ontology): the argument of

tradeoffs between storage and processing parsimony. Inferential processes are thought to be cognitively more complex, and in this sense less parsimonious, than simple disambiguation. Now, Bontly recalls that there is a standard response to this objection within the Gricean literature: specifically with regard to GCIs, it is claimed that “implicature calculation is largely unconscious and implicit: background assumptions can be taken for granted, steps can be skipped, and only rarely need the entire process breach the surface of consciousness” (Bontly, 2005: 307). As a consequence, it is held that implicatures of this sort are no more difficult than disambiguation. However, Bontly sees in this argument

an interesting difficulty that (to my knowledge) Griceans have never adequately addressed, for it is now quite unclear why this default interpretation should be considered an implicature rather than an additional sense of the expression. To say that it is a default interpretation is, after all, to say that speakers learn to associate that interpretation with the type of expression in question [...]. Short-circuited' implicature-calculation is thus hard to differentiate from disambiguation, making the Gricean hypothesis look more like a notional variant than a real competitor to the ambiguity hypothesis. (Bontly, 2005: 308)<sup>8</sup>

This said, Bontly attempts to get rid of the difficulty by observing that conventional meanings are inherently arbitrary while implicated meanings are not: “Familiarity with [a given] use [...] can remove the need to go through the canonical inference, but it doesn't change the fact that the use has a 'natural' (i.e., nonconventional) explanation. It doesn't change the fact that it is *calculable*” (Bontly, 2005: 308, his emphasis). In other words, the argument goes, although familiarity can cause implicatures to be short-circuited to the point that they become hard to differentiate from disambiguation, nonetheless short-circuited inferences are different from conventional meanings in that the former, but not the latter, are calculable. Therefore, Bontly suggests, short-circuited implicatures can ensure processing parsimony without becoming indistinguishable from conventional meanings.

But again, this move just amounts to making an appeal to what Bontly should prove: that if a meaning is *calculable* then one is justified to assume that it is not conventional. To put it differently, Bontly makes a good point when he observes that, from a cognitive point of view, short-circuiting of implicatures looks pretty much like the storage of associative relationships. If one nonetheless resists considering short-circuited implicatures as a sort of conventionalised meanings *for the reason that* the former (but not the latter) are calculable, then in practice she is making an appeal to MOR – to the principle that we should not posit new conventional meanings whenever they are calculable. The

---

8 But see note 4 above: what Bontly here calls “ambiguity hypothesis” should be rather called “polysemy hypothesis”.

circularity pops up again.

It could be argued that there is no circularity if, instead of making an appeal to MOR as an epistemological principle, Bontly were simply pointing at an objective, though dispositional, difference between mere conventional meanings and short-circuited implicatures. Bach (1998, which revisited Bach, 1995; see also Morgan, 1978) has seemingly made a similar point. As far as standardised implicatures are concerned, he says,

[t]he inference is compressed by precedent. But were there no such precedent, in which case a more elaborate inference would be required, there would still be enough contextual information available to the hearer for figuring out what is being conveyed. That is why special conventions are not needed for these special cases. (Bach, 1998: 713)

Bach seems to be suggesting here that short-circuited implicatures are distinguishable from genuine conventions not in terms of the actual processes involved, but instead in terms of a dispositional property: they could be substituted for by more elaborate inferences, in case it were required. However, as a matter of fact more elaborate inferences are normally not required to the extent that they are compressed by precedents. Therefore, from a cognitive perspective one may doubt that such a dispositional property has any import (in the normal case). Let me insist on this point. Bontly has identified a genuine difficulty in drawing a cognitive distinction between short-circuited implicatures and conventionalised interpretations. To the extent that they are compressed by precedents, inferences seem to be replaced by more straightforward associative processes, and intended meanings wind up being conventional. The fact that intended meanings could still be obtained by means of more elaborate inferences, were they not associatively recovered, does not change associative recovery into something different. As Mazzone (2013 a) puts it, “the mere fact that, when needed, a genuine inferential process may occasionally occur as a substitute for associative access does not change the latter into a half-inferential process”.

In sum, there seems to be a real difficulty for the notion of standardised implicatures: once it is recognised that implicatures can be standardised through repetition, then it is far from clear that they remain implicatures nonetheless. A wholly standardised implicature is presumably changed into a conventional meaning proper. A corollary of this observation is that, since conventionalisation is not an all-or-nothing affair, then we may reasonably find cases where meanings are just partially conventionalised. Phillips (2011) has provided an epistemological argument against MOR which is based precisely on this premise:

we can imagine a sorites series going from uses of expressions that are determinately not conventional (e.g. obvious implicatures), through those regularities in the borderline region that are neither determinately conventional nor determinately not conventional because the extra ingredients that are required for a convention to be in place are neither determinately present nor determinately not present, to those regularities that are clearly conventional. (Phillips, 2011: 6-7)

According to Phillips, the existence of a borderline region of indeterminacy casts some doubt on the overall epistemological strategy pursued by MOR. This strategy is in fact based on two assumptions: (1) that the linguistic data are indecisive between the implicature account and the polysemy account, and (2) that these two accounts are the best available explanation of the data. Only once these assumptions are made, can one appeal to some epistemological principle in order to favour one account over the other. However, as Phillips observes, while (1) is clearly true, (2) is probably false. There is in fact an alternative explanation of the data which is better than both the implicature and the polysemy account: if the data are indecisive between the two, given that conventionality comes in degrees, then the most plausible explanation is that the cases at issue are somewhat intermediate between being wholly conventional and being pragmatically calculated.

Not only does this line of thought pose a problem for MOR as an argument for the implicature account, it also calls our attention to a fact that cannot easily be accommodated by that account. As far as word uses are concerned, Griceans aim to give a pragmatic explanation – based on the notion of implicature – in any case except those for which everyone agrees that a conventionalisation has occurred. Moreover, in intermediate cases many Griceans are ready to apply the hybrid notion of standardised implicatures, that is, implicatures that are compressed by precedents. However, such a notion makes little sense from a cognitive point of view: as we saw above, it is disputable that a wholly short-circuited implicature is still an implicature “rather than an additional sense of the expression” (Bontly, 2005: 308). Therefore, the Gricean explanation for intermediate cases collapses into the polysemy account. But this difficulty in establishing a notion that is intermediate between full-fledged implicatures and fully conventionalised meanings also has consequences for the notion of implicature as such. Since conventionalisation comes in degrees (as emphasised by Phillips, 2011), one may wonder what happens to implicatures when implicated uses become more and more conventionalised. Do implicatures decline complementarily in a graded fashion? I know of no attempt to provide an answer to this issue, which is what I dubbed above “the continuum problem” for implicatures.

This is a problem, I maintain, for any theory of lexical pragmatics that is based on the notion of

implicature.<sup>9</sup> A sensible account of how word meanings are stored and recovered must plausibly assign some role to associative strength: word meanings can be more or less conventionalised as a function of regularity of use. As a consequence, if implicatures have to complement conventionalisation they should decline gradually as well. But then, we apparently need an account of implicature-like effects which does not imply any strong discontinuity between lexical retrieval and alleged implicatures. In the following section I will sketch the general lines of such an account.

#### 4. Lexical Pragmatics Without Inferences

As I said above, it is probably a main tenet of most current pragmatic theories inspired by Grice's work that implicatures have a strictly inferential nature, in the sense that they involve some sort of (possibly implicit) reasoning instead of mere associative recovery of coded information. Therefore, pragmatic processing is generally conceived of as an inferential affair, with the important exception of Recanati's (2004) account of lexical pragmatics. He proposes that the determination of explicit meaning, although being a pragmatic affair, is nonetheless ensured by wholly associative processes, to which he refers as primary pragmatic processes. By contrast, Recanati maintains that the transition from the explicit meaning of the utterance – fed by primary pragmatic processes – to its implicit meaning is ensured by secondary pragmatic processes, conceived in terms of genuinely inferential, personal-level processes. However, in Recanati's opinion primary pragmatic processes are apt to mimic genuine pragmatic inferences (Recanati, 2007). Therefore, in his view of lexical pragmatics the same associative mechanism is responsible for both lexical retrieval and implicature-like effects. In what follows, I am going to sketch a more explicit cognitive interpretation of this view (along the lines of Mazzone, 2011) than the one provided by Recanati himself.

For the sake of clarity, two main components must be distinguished in the account I propose. The first is the spreading of activation from linguistic expressions to associatively related representations within the conceptual network; the second component is what Recanati calls the “accessibility shift”,

---

<sup>9</sup> More generally, it is a problem for any theory of lexical pragmatics that is based on inferential mechanisms, possibly including explicatures in Relevance Theory. However, since Relevance Theory sensibly differs from the approaches considered so far – it conceives of pragmatic inferences in terms of specialized automatic processes, and it does not make any use of the notion of standardised implicature – I will address it separately in section 5.

occurring in the course of processing as soon as new information is made available to the cognitive system.

The first mechanism is essentially dependent on (provisionally) fixed relationships of accessibility, based on previously stored regularities. In practice, a linguistic expression can be associated with different strengths to a variety of concepts. More generally, any informational input may have a set of representations associated to it which are therefore accessible through spreading activation; the differential strength of those associations determines the order of accessibility for the representations in that set – that is, the order in which the activation spreads from a given informational item to the related representations. This temporal order of accessibility is correlated to how strongly each representation in the set is activated as a consequence of the activation of the informational item: representations that are accessed first are also (*ceteris paribus*) activated at the highest degree, since activation declines as it spreads.

The order of accessibility rooted in stored associations, however, cannot be the last word. Incoming information is usually processed within the context of both previous patterns of activation in the brain and immediately subsequent trains of activation due to further information. This may explain what Recanati calls “accessibility shift” in the course of processing. Representations that would have received the highest activation in the absence of any context can be outweighed by others in case supplementary information is provided. More generally, contextual information may enhance or reduce the activations induced by an isolated item, thus subverting the standard pattern of accessibility.

A crucial component in Recanati's account is the notion of schemata bridging informational items. A schema can be described as a representation tying two or more concepts together into a structured whole. In practice, any informational item may activate schemata which in turn spread the activation to their other components. Schemata are appealed to by Recanati in order to explain how accessibility shifts may promote the search for coherence in interpretation. In his account of primary pragmatic processes, the initial advantage of a given interpretation may be overridden due to the fact that a less accessible meaning has a better fit with – and therefore receives further activation by – some informational item via schemata of which they are both components. In this way, meanings that are more coherent with other informational items in the context come to be preferred over less coherent ones. There are two plausible assumptions underlying this account: first, that conceptual information is constituted of concepts multiply related to each other through schemata which specify the nature of their relationships;



and second, that schemata are actual components of the conceptual network, and therefore activation spread through them as well as through the concepts they tie together.

Why does Recanati hold that the processes described above are apt to mimic genuine inferential processes? The reason is that, once the role of accessibility shift and schemata is put into the picture, associative spreading activation is far from appearing as a random process. Quite on the contrary, it performs so as to obtain the same motivated results as if conclusions were drawn by way of rational, inferential processes. Let us consider the following example from Carston (2007):

- (6) Ann: I expected Jane to be here by now.  
Bob: She missed her coach.

The word form “coach” is potentially ambiguous between two possible meanings:  $COACH_1 = INSTRUCTOR$  and  $COACH_2 = BUS$ . Carston's (2007) account of how the addressee may come to conclude that “coach” refers to a bus is along the usual lines of RT. Specifically, she appeals to an inferential process where certain conclusions are drawn from certain assumptions, and in particular the crucial assumption is the following generalisation from experience:

- (7) MISSING A DESIGNATED  $COACH_2$  IS A REASON FOR A PERSON NOT ARRIVING WHEN EXPECTED.

However, this assumption is nothing but a schema in Recanati's sense, that is, a pattern connecting concepts through some specific kind of relation.<sup>10</sup> Here the relation at issue is one of causal explanation/rational justification: a person might not arrive when expected *for the reason that* s/he has missed the bus. The schema licenses the conclusion that presumably Bob means that Jane has missed a *bus*, since this interpretation makes his utterance contextually relevant by giving a reason for Jane's delay. But that conclusion would be equally prompted by an associative search for coherence driven by the schema at issue. Various word forms in Ann's and Bob's utterances (“expected”, “missed”, “coach”) may

---

10 Recanati (2004: 37) employs the example of the causal schema *STEAL (X) – IS ARRESTED (X)* in order to explain the interpretation of the pronoun in the sentence “John was arrested by a policeman yesterday; he had just stolen a wallet”. As to (7), it is conceptually equivalent to the causal schema *MISS THE COACH (X) – DO NOT ARRIVE WHEN EXPECTED (X)*.

be expected to spread activation to the schema insofar as the associated concepts (TO EXPECT, TO MISS, COACH<sub>2</sub>) are component of the schema, and this in turn would enhance the activation of the contextually relevant meaning of “coach”. In sum, what guarantees the conclusion is not the fact that the schema is employed in a genuinely inferential process, but rather the schematic information in itself, due to its power to constrain and drive the interpretation. It is such a schematic organisation of conceptual networks that makes it possible for associative processes to mimic genuine inferential ones.

I have argued for this view at more length in Mazzone (2011). What I want to emphasise here is that in this account there is no discontinuity between lexical retrieval and inferential effects. Schemata can constrain utterance interpretation by way of the very same mechanism involved in lexical retrieval, that is, the dynamics of activation they take part in.<sup>11</sup>

To be sure, in the above example we assumed that the intended meaning is already conventionalised, and therefore encoded by the speakers involved: it is one of the lexical meanings of the ambiguous word form “coach”. However, as we observed in the third section, ambiguity (and polysemy

---

<sup>11</sup> An anonymous referee has urged that I should make clear what (how much) is encompassed by the notion of schema, insofar as my account places huge weight on it. The concern is that the notion cannot be broad enough to cover the entire contribution of encyclopedic/general knowledge to utterance understanding. Furthermore, the same referee wonders whether my claim of a continuity between lexical retrieval and inferential effects is compatible with experimental work such as that of Rubio Fernandez (2007) which seems to show that meaning selection and meaning construction are different processes. I cannot fully answer these questions here. Let me just say this much. As to the first issue, I adopt a very general notion of schema, according to which any pattern of regular covariation between two or more concepts is a schema. I believe that this notion is justified by its role in explaining how concepts are both formed (Mazzone, in preparation) and used in cognitive processing (Mazzone, in press). As to the other issue, I do not claim that there is no discontinuity at all between meaning selection and meaning construction. Specifically, I do not deny that meaning construction might involve working memory and attentional processes. On the contrary, in Mazzone (2013 b, c) I explicitly maintain that, for a full explanation of utterance understanding, associative processes need to be complemented by attentional resources. However, this must be compatible with the fact underlying the continuum problem, that is, the fact that pragmatic inferences must decline progressively to the extent that uses become more and more conventionalised. This is easily explained if we assume that the same associative dynamic is involved in the two cases: the indirect associative relation between two (or more) concepts may be gradually short-circuited by repetition, and this in turn affects how those concepts may activate each other – (more or less) indirectly, via schemata, and (more or less) directly. Incidentally, the opposition between continuity and discontinuity in cognitive processing might not be as neat as it is often assumed (see also the end of section 5; more on this in Mazzone, 2013 c.).

as well) does not dispense with inferences or inference-like effects, since, in order for the appropriate meaning to be selected, the context needs to be taken into account anyway, and schematic assumptions are needed to bridge intended meaning and contextual items.

A different case is when the intended meaning is not encoded by the hearer. In this case, the recovery of the meaning is more indirect. Let us suppose, for instance, that the listener has not (yet) encoded the concept BUS as a conventional meaning of the word form “coach”. We might nonetheless assume that her conceptual system contains the information expressed by the schema (7), although not in that precise form since by hypothesis that system does not include a concept COACH<sub>2</sub> (= BUS) associated to the word form “coach”. Then the schema would rather be as follows:

(8) MISSING A DESIGNATED BUS IS A REASON FOR A PERSON NOT ARRIVING WHEN EXPECTED

Even in this case, the schema could help the listener to identify BUS as the intended meaning of “coach”, though in a more indirect way. The schema could be weakly activated – together with other schemata providing alternative possible reasons why a person could not arrive when expected – by Ann's utterance (“I expected Jane to be here by now”). Further, the word form “missed” in Bob's answer would probably add activation to a subset of those schemata, specifically the ones involving missing buses, trains etc. Then, the context (or stored knowledge) might provide further clues leading the listener to select MISSING THE BUS amongst the possible reasons. In sum, it is not inconsistent to think that, insofar as the intended meaning is recovered at all<sup>12</sup>, this could be ensured again by the activation of schemata bridging contextual clues (or stored knowledge) and the intended concept.

Besides, in a number of cases it happens that the lexicalised meaning of a word form and the meaning attributed to it in a circumstance are tied by some sort of relation (metonymy, similarity etc.). In those cases, although the concept which is provided by the schema is not itself a lexicalised meaning of

---

12 An anonymous referee has suggested that in this case confusion or misinterpretation, and not the recovery of the word meaning, is the most likely outcome: “the hearer would try to find some way of making sense of the proposition that Jane missed her instructor and that this has somehow delayed her arrival”. I do not argue with this suggestion. My aim is just to sketch the general lines of how even the most difficult case – that is, the case in which the lexical meaning is unknown to the hearer – can be accounted for in terms of activation of schemata. But this does not commit me to the claim that the hearer necessarily succeeds in recovering the meaning.

the word form, there is presumably some further schema (be it metonymic, metaphoric, etc.) connecting the lexical meaning and the intended one. These are clear cases in which the intended meaning is calculable, that is, it can be derived from the lexicalised one in a motivated manner, and in which polysemy (versus homonymy) may result when the intended meaning becomes lexicalised.

The crucial point is that the proposed account allows for a gradual transition between the cases considered above, due to the fact that lexical retrieval and the contribution of schemata are just different facets of the same activation dynamics. Thus, when the intended meaning is not lexicalised so that some wrong concept (or no concept at all) tends to be activated by simple lexical retrieval, associative activation of schemata may be sufficient to yield the required accessibility shift and drive the search for the right concept. In the end, the intended meaning may become progressively conventionalised: this equates to the formation and strengthening of a direct connection between the word form and the meaning, so that the more indirect route mediated by other schemata is short-circuited and eventually bypassed. Since in this view the indirect route does not involve genuine inferences but rather associative mechanisms mimicking inferences, the gradual decline of this route as a consequence of conventionalisation does not raise any difficulty.

It should be noticed that this approach is not incompatible with the intuition according to which, at least in some cases, word uses that are calculable might not be (much) conventionalised in spite of their frequency. The point is that the frequency with which a word form conveys a meaning is just one factor of its conventionalisation. Other plausible factors are the frequency of previously lexicalised meanings, the absence of synonyms, the independence of new uses from old meanings, the determinacy of word uses, the grammatical category – and possibly others. Let us consider, for instance, the metaphorical sentence

(9) Juliet is the sun.

Could the phrase “the sun” come to be lexically associated to – let us say – the concept BEAUTIFUL, even in case it were frequently used in that sense? The answer must be mostly negative, for in this case almost all the factors cited above disfavour lexicalisation. The word form “sun” is basically used to express the concept SUN (as to the factor *frequency of previously lexicalised meanings*); the concept BEAUTIFUL is already lexicalised in English (*presence of synonyms*); the phrase “the sun” in contexts such as (9) has a much more indeterminate meaning than BEAUTIFUL (*determinacy of word uses*); the

phrase “the sun” belongs to a different grammatical category than the meaning BEAUTIFUL would require.

For another example, let us consider the metonymic sentence (discussed in Nunberg, 1977):

(10) The ham-sandwich left without paying.

The phrase “the ham-sandwich” is intended as referring to the person who ordered the sandwich, thanks to a transfer from one component to another of a café/restaurant schema. Now, it is not only that the literal meaning is no less (and probably more) frequent than the figural one (*frequency of previously lexicalised meanings*); moreover, this case is so construed that the literal meaning is salient in the context and presumably has a role in retrieving the figural meaning (there is little *independence of new uses from old meanings*). Frequency presumably favours lexicalisation of secondary meanings in inverse proportion to how much the contexts of use evoke primary meanings.

In sum, in both the examples above – (9) and (10) – one should not expect a significant amount of conventionalisation. But this is not evidence that whenever uses are calculable then they are calculated (by way of genuine inferences) instead of conventionalised. As a matter of fact, uses that are calculable are nonetheless conventionalised in many other cases: we have mentioned dead metaphors, but this also seems to apply to “dead metonymies” such as “White House” used to refer to the President of the United States, and so on and so forth. Whether calculable meanings are conventionalised or not depends on frequency of word use together with the interacting factors we considered above.

Given that, by this account, various factors of conventionalisation – including the frequency with which a word form conveys a given meaning and the frequency of previously lexicalised meanings – have continuous values, the overall picture is still compatible with the idea that there is a continuum between wholly conventionalised and non-conventionalised uses and that, consequently, a gradualist account is needed at the cognitive level.<sup>13</sup>

## 5. Relevance Theory and the Continuum Problem

---

13 Remember that strictly speaking I am not concerned here with conventions *per se*, but with how conventions are implemented in individual cognition.

In view of our above considerations, it is interesting to briefly consider relevance theorists' position in lexical pragmatics. Insofar as they too insist on the genuinely inferential nature of pragmatic processing, their position might seem to be plagued by the continuum problem just as much as the more traditional Gricean accounts. Some aspects of their approach, however, suggest the possibility of a different interpretation, which is more in line with my proposal in the previous section.

To start with, Wilson and Carston (2006, 2007) have recognised that associative processes may have a role to play in lexical pragmatics (more on this point in Mazzone, 2011: section 3).<sup>14</sup> Associations would determine the initial accessibility of concepts and, therefore, the initial probability that a concept is a component of the explicit meaning of utterances, although this meaning is eventually adjudicated by the properly inferential component of the process. Explicit meaning and contextual assumptions are in fact held to be the premises from which contextual implications (implicated conclusions) have to be inferred. Relevance theorists, however, have a non-linear conception of this inferential process. They speak of a mutual adjustment mechanism working in parallel without a rigid sequential order. This means that the premises of the inference might be recovered, or at least adjusted, with the help of backward inferences from the conclusion, in case the conclusion is made independently available thanks to contextual clues. Specifically, the hypotheses about explicit content (and the concepts that compose it) can be revised as a function of both the contextual assumptions and the conclusion(s) involved.

As an example consider the following exchange (from Wilson and Carston, 2006: 422), where the first speaker asks the question (11a) about the addressee's younger sister Caroline, who is manifestly not a princess, and then the addressee answers as in (11b):

- (11) a. Will Caroline help us clear up the flood damage?  
b. Caroline is a princess.

The meaning of “princess” in (11b) has to be modulated in accordance with the context so as to convey the concept PRINCESS\*, which differs from the lexicalised concept PRINCESS. Specifically, Wilson and

---

14 For instance, Wilson and Carston (2006: 429) claim that “associative links may affect the outcome of the mutual adjustment process by altering the accessibility of contextual assumptions and implications, but the resulting overall interpretation will only be accepted as the speaker’s intended meaning if it satisfies the hearer’s expectations of relevance and is properly warranted by the inferential comprehension heuristic discussed [above in that paper]”.

Carston propose that the following inference is drawn:

- (12) a. **Explicit content:** CAROLINE IS A PRINCESS\*
- b. **Contextual assumption:** A PRINCESS\* DOESN'T CLEAR UP FLOOD DAMAGE
- c. **Contextual implication:** CAROLINE<sub>x</sub> WON'T HELP US CLEAR THE FLOOD DAMAGE

Wilson and Carston observe that the contextual assumption in (12b) is unlikely to be stored in the encyclopedic entry for PRINCESS, but they suggest that “this feature should be straightforwardly derivable in the course of the mutual adjustment process, by a combination of forward inference from existing encyclopaedic features (e.g. UNUSED TO PERFORMING MENIAL TASKS, UNACCUSTOMED TO MANUAL LABOUR), and backward inference based on the expected type of conclusion in [(12C)]” (Wilson and Carston, 2006: 422). Clearly, the inferential process is not taken here to occur exclusively in a forward direction from the premises to the conclusion. On the contrary, there is a specific expectation about how the utterance in (11b) is likely to achieve relevance as an answer to (11a) (cfr. Wilson, 2004: 354), and therefore there is some expectation about the most likely contextual conclusion. This conclusion, in turn, contributes to the recovery of the premises – specifically, the contextual assumption (12b) and the concept PRINCESS\* in the explicit content (12a).

Such a mutual adjustment is exactly what one would expect in case the process at issue were associative. In associative processes different representations may act on each other as mutual constraints, independently from their order of activation. Then one would be tempted to conceive of the described process as inferential in the limited sense that there are schemata connecting the involved representations so as to potentially form a justificatory structure – that is, a structure that could be appealed to in order to inferentially justify the conclusion – although schemata and representations act on each other as associative constraints, not literally as sequential steps of an inferential derivation. In this case, we could claim that associative processes mimic inferential ones, just as Recanati (2007) suggested. An interpretation of RT along these lines seems consistent with the following claim made by Wilson and Carston (2007: 244): “From a cognitive point of view, all inferential relationships are also associations: an inferential mechanism establishes systematic correspondences between (constituents of) premises and (constituents of) conclusions” (see also Sperber and Wilson, 2008).

In sum, one could weaken the thesis of genuinely inferential processes to the point of saying that there are inferences provided that justificatory structures are involved, whatever the cognitive mechanisms by which this occurs. Of course, I would agree that pragmatic processes are inferential in this specific sense. However, relevance theorists would hardly be content with the idea that inferential pragmatic processes can be accounted for in wholly associative terms. While they clearly concede that the associative dynamic of activations is responsible for feeding concepts into the explicature (and also for feeding contextual assumptions into the derivation), they seem to assume that the inferential stage of processing requires a cognitive process of a quite different sort than the associative one. In my opinion, though, this assumption of a cognitive discontinuity might not be easy to reconcile with the RT's framework for lexical pragmatics described above, for three reasons.

First, within that framework all the inferential steps need to be open to the operations of associative processes. If inferences were thought of as sequential derivations based on a set of rules, then the conclusions could be thought to simply follow once the premises are set. But given the hypothesis of a mutual adjustment process, even conclusions must be associatively sensitive to contextual inputs in much the same way as explicit content and contextual assumptions are. Only if conclusions are independently activated by the context can they license backward inferences driven by specific expectations of relevance. Moreover, not only can any of the inferential steps be (provisionally) determined by contextual items by means of associative activation, it also goes the other way around: whenever mutual adjustments in the inferential steps occur, then the weight (i.e., the activation) of the related contextual items can be expected to change in turn. In other words, it is far from clear that associative processes can be restricted to a one directional flow that stops when explicature and contextual assumptions are fed into the derivation.

Second, the assumption of a mutual adjustment makes the picture more complex also in the sense that some sort of monitoring seems to be required. In a standard conception of inferential processes, derivations can be thought of as blind applications of rules. But in RT's framework, the mechanism needs to read – so to speak – the changes intervening in subsequent inferential steps in order to change accordingly the previous steps of the derivation. An associative account has no difficulties in coping with this issue, since changes in the activation of any representation automatically affect the activation of others provided that there exist schemata associatively bridging the two. On the other hand, RT has never specified, to my knowledge, how this monitoring and adjusting would be accomplished: it is one thing to



say that a mutual adjustment between premises and conclusions has to occur, it is quite another to provide an explicit cognitive explanation of how this can be done. In this sense, either the associative dynamic is allowed to extend beyond the determination of premises and conclusions and to accomplish the assessment of the inferential structure as well, or else RT needs to provide a different account of how the monitoring-and-adjusting process is accomplished.

Third, it is in fact not clear why one should presume that the surrounding associative dynamics has to stop on the boundary of the inferential component. Let us recall, for example, the discussion of (6), and in particular the schema invoked by Carston (2007) in order to arrive at the right interpretation of “She missed her coach”:

(7) MISSING A DESIGNATED COACH<sub>2</sub> IS A REASON FOR A PERSON NOT ARRIVING WHEN EXPECTED.

Once we admit that such a schema is part of the subjects' stored knowledge, there is no reason to assume that it is not involved in the general associative dynamic of our mind: not only can the schema be activated by contextual items, it also can activate in turn other representations counting as (components of) premises or conclusions within justificatory structures. Recanati (2004) has proposed that schemata such as (7) add activation to any item that happen to fit them. It is not difficult to imagine that activation may similarly propagate from one step to another of justificatory structures (i.e., premises and conclusions), either because they share some concepts or because there are abstract schemata connecting the premises and conclusions involved. Generally speaking, there is no reason to think that the associative dynamic has to stop at the point at which assumptions are fed to the alleged inferential component.

These considerations seem to show that, in accounting for mutual adjustment in pragmatic derivations, it is not easy to draw a demarcation between the associative stage and an alleged inferential process that would differ in kind. Specifically, it is not clear why the associative dynamic should stop when premises and conclusions are set (first and third point) and, if this were the case, how then the monitoring-and-adjusting step would be accomplished (second point). RT seems to assume that, while premises (and possibly conclusions) of pragmatic inferences are accessed through a merely associative route, they are evaluated by a different cognitive process. Contrary to this hypothesis, my previous considerations aim to show that the same dynamic of activation mediated by schemata can explain both

how premises and conclusion(s) of pragmatic inferences are accessed and how premises and conclusion(s) are adjusted to each other and to the context in a motivated manner.

In sum, RT is clearly committed to a quite non-standard account of inferential processing in pragmatics, according to which inferences are automatic and based on parallel mutual adjustment. This view does not seem prone to the continuum problem affecting the more standard Gricean accounts, and it seems naturally accommodated within a thoroughly associative view. However, relevance theorists seem to insist instead on a cognitive discontinuity between the associative and the inferential component of the process. I have argued, though, that the thesis of a cognitive discontinuity raises some problems. First, there is no reason to assume that the associative dynamic has no resources to ensure the rational adjustment of premises and conclusion(s) to each other and to the context. And second, if this were the case, RT would be in need to specify which mechanism can ensure such adjustment.

All this is not to say that associative processes are all there is to pragmatic processing. On the contrary, I have recently proposed that in language comprehension pure spreading activation must be complemented by working memory and attentional processes (Mazzone, 2013 b, 2013 c). However, although I cannot discuss the issue in detail here, the distinction between automatic and attentional processes can hardly correspond to the distinction between forming interpretive hypotheses and confirming/disconfirming them. In the first place, attentional processes can be thought to differ from automatic associative ones essentially in terms of the scale of activation/inhibition in the brain. As a matter of fact, the general dynamic of activation and inhibition in the brain takes place at different scales, spanning from most local to most global, with (most) global processes implementing executive functions and specifically attentional processing. Thus, there is no discontinuity in the strict sense between associative and attentional processes: they are part of the same dynamic of activation/inhibition just at different scales. In the second place, attentional processes do not provide new information or specialized rules in order for rational decisions to be made. They change instead the overall dynamic of activation from local to global – and from parallel to serial – so that only part of the information associatively available to the system for the purpose of rational assessment is specifically focussed and assigned a role in actively driving cognition.

Therefore, although in the normal case conscious attention is involved in utterance understanding, it would not be correct to say that in principle attention is necessary in order for interpretive hypotheses to

be confirmed/disconfirmed. On the contrary, automatic associative processes have in principle the resources to accomplish the mutual adjustment between premises and conclusions in a motivated manner.

## **6. Conclusions**

In this paper, I have focussed on a problem for Grice's account of lexical pragmatics, which I have referred to as “the continuum problem”, and I have proposed a solution to it.

The existence of a continuum between lexicalised meanings and alleged implicatures raises a problem for both Modified Occam's Razor and the very notion of implicature, at least as it is used in lexical pragmatics. The argument according to which in case of uncertainty between conventionalisation and implicatures we should prefer an account based on the latter loses its grip once we recognize that there can be intermediate cases between pure lexicalised meanings and pure implicature-like cases. Furthermore, those intermediate cases raise a problem for a cognitive interpretation of the notion of implicature since, while it is easily agreed that conventionalisation may come in degrees, it is far from clear that the processes delivering implicatures may decline accordingly.

A possible solution to this problem has been found in an associative account, where implicature-like effects are explained as the result of schemata which constrain the selection of contextually appropriate interpretations.

## **Acknowledgments**

A previous draft of this paper has been discussed in a seminar at University College London in May 2013. I want to thank all the participants – Barbara Eizaga Rebolgar, Alison Hall, Xiaohong Jiang, Diana Mazzarella, Isabelle Needham-Didsbury, and especially Robyn Carston and Deirdre Wilson – for their precious comments. I have also to thank the anonymous referees for their insightful comments.

## References

- Bach, Kent. 1995. Standardization vs. Conventionalization. *Linguistics and Philosophy* 18: 677-686.
- Bach, Kent. 1998. Standardization revisited. In A. Kasher (ed.), *Pragmatics: Critical concepts. Vol IV*, 712-22. London and New York: Routledge.
- Barsalou, Lawrence. 1992. *Cognitive Psychology: An Overview for Cognitive Scientists*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Bontly, Thomas D. 2005. Modified Occam's Razor: Parsimony, Pragmatics, and the Acquisition of Word Meaning. *Mind & Language* 20 (3): 288-312.
- Carston, Robyn. 2007. How many pragmatic systems are there? In M.J. Frapolli (ed.), *Saying, Meaning, Referring. Essays on the Philosophy of François Recanati*, 1-17. New York: Palgrave.
- Carston, Robyn. 2010. Explicit Communication and 'Free' Pragmatic Enrichment. In B.Soria and E. Romero (eds.), *Explicit Communication: Robyn Carston's Pragmatics*, 217-87. New York: Palgrave.
- Croft, William and David Alan Cruse. 2004. *Cognitive Linguistics*. Cambridge: Cambridge University Press.
- Dehaene, Stanislas, Jean Pierre Changeux, Lionel Naccache, Jerome Sackur and Claire Sergent. 2006. Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, 10(5): 204–11.
- Devitt, Michael. 2004. The case for referential descriptions. In M. Reimer & A. Bezuidenhout (eds.), *Descriptions and Beyond*, 280-305. Oxford: Oxford University Press.
- Devitt, Michael. 2007. Referential Descriptions and Conversational Implicatures. *European Journal of Analytic Philosophy*, 3: 7-32.
- Gazdar, Gerald. 1979. *Pragmatics. Implicature, Presupposition, and Logical Form*. New York: Academic Press.
- Goldberg, Adele. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago: University of Chicago Press.
- Grice, Herbert Paul. 1978. Further Notes on Logic and Conversation. In P. Cole & J. L. Morgan (eds.), *Syntax and Semantics, vol. 9, Pragmatics*, 183-97. New York: Academic Press.
- Grice, Herbert Paul. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University

Press.

Horn, Laurence R. 1972. *On the Semantic Properties of Logical Operators in English*.

Bloomington: IULC.

Horn, Laurence R. 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In D. Schiffrin (ed.), *Meaning, Form, and Use in Context*, 11-42. Washington: Georgetown University Press.

Horn, Laurence R. and Samuel Bayer. 1984. Short-circuited implicature: A negative contribution. *Linguistics and Philosophy*, 7: 397-414.

Mazzocco, Michele. 1997. Children's interpretations of homonyms: A developmental study. *Journal of Child Language*, 24: 441-67.

Mazzone, M. 2011: Schemata and associative processes in pragmatics. *Journal of Pragmatics*, 43, 2148–2159.

Mazzone, M. 2013 a: A pragmatic Pandora's box. Regularities and defaults in pragmatics. In F. Liedtke and C. Schulze (eds.), *Beyond Words. Content, Context, and Inference*. Berlin: Mouton de Gruyter.

Mazzone, M. 2013 b. Attention to the speaker. The conscious assessment of utterance interpretations in working memory. *Language & Communication*, 33(2), 106-114.

Mazzone, M. 2013 c: Automatic and controlled processes in pragmatics. In A. Capone, F. Lo Piparo, M. Carapezza (eds.), *Perspectives on Linguistic Pragmatics*. Berlin: Springer.

Mazzone, M. in press, Crossing the associative/inferential divide: ad hoc concepts and the inferential power of schemata, *Review of Philosophy and Psychology*.

Mazzone, M in preparation, How concepts are formed. The relationship between schemata, theories, and similarity

Morgan, Jerry. 1978. Two types of convention in indirect speech acts. In P. Cole (ed.), *Syntax and Semantics, vol. 9, Pragmatics*, 261–280. New York: Academic Press.

Nunberg, Geoffrey D. 1977. *The pragmatics of reference*. PhD dissertation, City University of New York.

Phillips, Ben. 2012. Modified Occam's Razor. *Australasian Journal of Philosophy*, 90 (2): 371-381.

- Recanati, François. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.
- Recanati, François. 2007. Reply to Carston. In M. J. Frapolli (ed.), *Saying, Meaning, Referring. Essays on the Philosophy of François Recanati*, 49–54. New York: Palgrave.
- Reimer, Marga. 1998. Donnellan's distinction/Kripke's test. *Analysis*, 58: 89–100.
- Rubio Fernandez, Paula. 2007. Suppression in metaphor interpretation: Differences between meaning selection and meaning construction. *Journal of Semantics*, 24: 345-71.
- Sadock, Jerrold. 1978. On testing for conversational implicature. In P. Cole (ed.), *Syntax and Semantics, vol. 9, Pragmatics*, 281–297. New York: Academic Press.
- Searle, John. 1975. Indirect speech acts. In P. Cole and J. Morgan (eds.), *Syntax and Semantics 3: Speech Acts*, 59-82. New York: Academic.
- Sober, Elliott. 1988. *Reconstructing the Past*. Cambridge, MA: MIT Press.
- Sperber, Dan. 2005. Modularity and relevance: How can a massively modular mind be flexible and context-sensitive? In P. Carruthers, S. Laurence and S. Stich (eds.), *The innate mind: Structure and content*, 53–68. Oxford, England: Oxford University Press.
- Sperber, D. & Wilson, D. (2008). A deflationary account of metaphors. (In R. W. Gibbs (ed.), *The Cambridge handbook of metaphor and thought* (pp.84-105). Cambridge: Cambridge University Press.)
- Walker, Ralph. 1975. Conversational implicatures. In S. Blackburn (ed.), *Meaning, Reference, and Necessity*, 133–81. Cambridge: Cambridge University Press.
- Wilson, Deirdre, 2004. Relevance and lexical pragmatics. *UCL Working Papers in Linguistics*, 16: 343-360.
- Wilson, Deirdre and Robyn Carston. 2006. Metaphor, relevance and the emergent property issue. *Mind and Language*, 21 (3): 404–433.
- Wilson, Deirdre and Robyn Carston. 2007. A unitary approach to lexical pragmatics: relevance, inference and ad hoc concepts. In N. Burton-Roberts (ed.), *Pragmatics*, 230–260. New York: Palgrave.