

# The Paradox of Asserting Clarity

Chris Barker and Gina Taranto

UC San Diego

## 1. The Dilemma

Standard wisdom (Stalnaker 1979:325, van der Sandt 1992:367, etc.) holds that assertions are felicitous only if they add new information to the common ground. After all, what use could it be to claim that a proposition is true if it is already accepted as true? In this paper we suggest that this question is not rhetorical. Our answer is that some sentences can have side effects besides adding information to the common ground, and that sometimes it is worth asserting a sentence entirely for the sake of its side-effects.

To motivate our claim, consider a variation on Partee's marble example:

- (1) a. Exactly two out of three marbles are on the table.
- b. One marble is not on the table.
- c. It's under the couch.

(1b) is entailed by (1a); it adds no new information about the situation under discussion. However, it causes the creation of a discourse referent for the missing marble, which allows the pronominal reference in (1c). Without (1b) it would be infelicitous to use the pronoun in (1c). Thus, as pointed out in Beaver (2002:172), it is possible to assert a sentence purely for the sake of its side-effects, here, building a discourse referent to facilitate anaphora.

This paper presents a case-study of the semantics of *clear*, which we take to be a Discourse Adjective following Taranto (2002a, 2002b). The central example we consider is given in (2).

- (2) It is clear that Mary is a doctor.

Intuitively, for (2) to be true, the discourse participants must each possess all the knowledge they need to conclude that Mary is a doctor before (2) is uttered. If either is not already convinced that Mary is a doctor, then the proposition isn't

clear at all. But if it is already evident that Mary is a doctor, then asserting (2) adds no new information to the common ground, suggesting the puzzle in (3).

(3) Lemma 1: If (2) is true, it adds no new information to the context, so why bother to assert it?

To begin, note that it is not appropriate to utter (2) if the fact that Mary is a doctor has just been asserted. Thus the discourse in (4) is distinctly odd.

(4) a. I just learned that Mary is a doctor.  
b. #Clearly, Mary is a doctor.

Whatever uttering (4b) is supposed to accomplish, it can't be done immediately after uttering (4a). For (4b) to be felicitous, the fact that Mary is a doctor cannot be part of the common ground. As soon as Mary's doctorhood is in the common ground, as would be the case immediately after an utterance of (4a), it becomes impossible to assert that it is clear that she is a doctor. This suggests that perhaps an utterance of (2) adds new information after all, namely the information that Mary must be a doctor. But this gives us a second puzzle:

(5) Lemma 2: Assume (2) entails Mary is a doctor. If the speaker has decided that Mary is a doctor on the basis of information that the hearer also has, then (2) only adds information if the speaker assumes that the hearer has not come to the same conclusion- in which case it is not in fact clear that Mary is a doctor, otherwise the hearer would have realized it on her own.

The second lemma might be rephrased as "how can (2) be uttered without assuming the hearer is an idiot?" Perhaps it can only be used in situations in which the hearer has all the evidence he needs to realize that Mary is a doctor, but fails to take that last logical step. Consider the propositions in (6).

(6) a. Mary is holding a stethoscope.  
b. Mary is wearing a lab coat.  
c. Mary knows lots of Latin morphology.

Assume the propositions in (6) are true, and that they are in the common ground by virtue of having just been uttered by B. At this point A replies with her version of (3): "Clearly, you dolt, Mary is a doctor."

Resolving our dilemma requires developing a specific model of context update, and elaborating that model in order to handle vagueness.



Now consider asserting (12a) in a situation in which the hearer does not know either Pat's gender or Pat's marital status:

(13)Initial Context C:	Pat is a man	Pat is a woman
	Pat is married	{ $w_1, w_2$ $w_3, w_4,$
	Pat is not married	$w_5, w_6,$ $w_7, w_8$ }

After an utterance of (12a), only worlds in which Pat is an unmarried man, namely, worlds  $w_5$  and  $w_6$ , survive update. The updated context is strictly smaller than the initial context, so asserting (12a) has added new information.

Now consider an assertion of (12b) following (12a). This does not eliminate any worlds, since all of the worlds in which Pat might have been a woman were already eliminated after the utterance of (12a). Thus asserting (12b) following an assertion of (12a) violates (11). The analysis correctly predicts that it would be infelicitous to utter (12a) immediately followed by (12b): *#Pat is a bachelor, and Pat is also a man*. Thus, in general, the Stalnakerian model guarantees that when the information expressed by a sentence is added to a context, the information corresponding to all of the entailments of that sentence are also added. Since we are committed to the Stalnakerian model, we must reject the missing entailment hypothesis.

Since people are not always consistent (i.e., they are capable of simultaneously believing a proposition and denying its consequences), we might conclude that contexts can be similarly inconsistent. This would mean that, in this respect, the Stalnaker model of context update is inaccurate, and this is a flaw in the model. But another alternative that we should consider first is that the facts in (6) do not actually entail that Mary is a doctor. This alternative respects the fact that there are a number of possibilities for how the world might be, and that these possibilities are ordered in terms of their plausibility:

- (14)Possible explanations for the evidence suggesting Mary is a doctor, from most likely to least likely:
- $w_9$  = Mary is a doctor.
  - $w_{10}$  = Mary is a doctor, though she learned her Latin in high school.
  - $w_{11}$  = Mary is not a doctor, she's getting ready for a Halloween party.
  - $w_{12}$  = The stethoscope is Mary's brother's, but Mary is a doctor too.
  - $w_{13}$  = Mary is not a doctor, but the CIA wants us to believe that she is.

Each scenario in (14) corresponds to one way the world might be, and nothing in the sentences in (6) rules out any of these possibilities. Because in some of these worlds Mary isn't a doctor, update with (3) will eliminate those worlds in

which Mary isn't a doctor. However, asserting *Mary is a doctor* will achieve the same result:

(15)  $\{ w_9, w_{10}, w_{11}, w_{12}, w_{13} \} + \textit{Mary is a doctor} = \{ w_9, w_{10}, w_{12} \}$

We are now faced with the following question: why not just assert that Mary is a doctor? Why ever assert it is *clear* that Mary is a doctor? We claim that a speaker might be reluctant to assert that Mary is a doctor precisely because Mary might not be a doctor. There are other possibilities that are still live, and we know from Grice that it would be uncooperative to claim that Mary is a doctor without being absolutely sure.

If this is on the right track, (2) might be used to signal that a speaker doesn't have enough information to flatly assert that Mary is a doctor. That is, clarity is asserted only in contexts in which there is some lingering uncertainty that the complement is in fact true. But if this is right, it is extremely peculiar, since it means that we have reconstructed our original paradox, only in reverse:

(16) The reconstructed paradox: *It is clear that p* is asserted only in situations in which it is in fact **not** clear that *p*!

We believe the key to resolving this paradox lies in characterizing how the grammar deals with degrees of probability. The appropriateness of asserting clarity depends on degrees of probability of different explanations for the facts.

Situations in which the applicability of a predicate depends on degrees are well known in the literature of vagueness (Fine 1975, Williamson 1994, Kennedy 1997). We will argue, however, that *clear* is not an ordinary vague predicate.

### 3. Vagueness

Vagueness is about where to draw the line between having or not having a property. A predicate like *tall* is vague because in a given situation, it often isn't clear exactly how tall you need to be to count as tall. Assume that in any given discourse situation there is a standard for how tall a person needs to be in order to count as tall. Following Barker (2002), we write this as:

(17)  $\mathbf{d}(c)(\llbracket \textit{tall} \rrbracket)$

Here  $\mathbf{d}$  is a delineation function (Lewis 1970) which takes a situation  $c$  and an adjective meaning and returns the vague standard for that adjective in the given situation. Then we can characterize the truth conditions of (18a) as (18b).

- (18)a. *Bill is tall.*  
 b. The maximal degree to which Bill is tall is at least as great as  $\mathbf{d}(c)(\llbracket tall \rrbracket)$ .

(18a) can be used either descriptively or metalinguistically (Kyburg and Morreau 2000, Barker 2002). The simplest way is descriptively. Assume a situation  $c$  in which the standard for human tallness is exactly six feet. In  $c$ , the delineation function applied to the adjective *tall* returns the vague standard of 6 feet, as in (19).

$$(19) \mathbf{d}(c)(\llbracket tall \rrbracket) = 6'0''$$

The dialogue in (20) illustrates the descriptive use of a vague adjective.

- (20)a. What is Bill like?  
 b. Bill is tall.  
 c. The maximal degree to which Bill is tall is at least as great as 6'0''.

Relying on our knowledge about the local standard for tallness, the interlocutors have learned a lower bound on Bill's height. The assertion of *Bill is tall* has added descriptive information about the way the world is.

To illustrate the metalinguistic use, imagine a speaker and hearer both know a lot about Bill, including the exact degree to which Bill is tall, which is 6'1''. In contrast with our previous scenario, however, the standard for human tallness is more obscure. The interlocutors have their individual ideas of how tall one has to be to count as tall, but they don't know if their individual standards coincide with their interlocutor's standard. They might proceed as in (21).

- (21)a. What counts as tall around here?  
 b. Well, see Bill over there? Bill is tall.  
 c. The maximal degree to which Bill is tall is at least  $\mathbf{d}(c)(\llbracket tall \rrbracket)$   
 d.  $\mathbf{d}(c)(\llbracket tall \rrbracket) \leq 6'1''$

In this situation, an assertion of *Bill is tall* provides no new information about Bill, since the discourse participants knew exactly how tall Bill was to begin with. They do however gain information about the prevailing standard for tallness: it must be less than Bill's height.

We claim that when a speaker asserts and a hearer accepts a claim that *Bill is tall*, they reach a tacit agreement about the contextually relevant constraint on tallness. That is, they take a concrete step towards synchronizing their

individual standards for tallness, and they can rely on this in future discourse. This is a metalinguistic, rather than a descriptive use.

These two aspects of meaning can be easily modeled building on Stalnaker's (1998) notion of context update. We need only adopt his natural assumption that during a conversation, some things are certain about the world: a conversation is taking place, the speaker is speaking, the hearer is being addressed, and so on. Thus, every possible world in the initial context will be a world in which the conversation underway is taking place. Following this, we conclude that one way in which worlds may vary is in the value of the delineation function for the version of the conversation in that world. Uses of sentences involving vague predicates are not necessarily purely descriptive or purely metalinguistic— they are usually a mixture of both. That is, discourse improves mutual knowledge both concerning the world under discussion as well as concerning the nature of the discourse itself. This is not surprising, of course, since the discourse itself is part of the world, and therefore a legitimate target for reducing ignorance.

#### 4. Analysis of *clear*

Besides raw intuition, the vagueness of *clear* is easy to prove, since it is possible to explicitly talk about the degree to which a proposition is clear.

- (22)a. It is becoming clear that Mary is a doctor.
- b. It is reasonably clear that Mary is a doctor.
- c. It is very clear that Mary is a doctor.
- d. It is painfully clear that Mary is a doctor.

Our preliminary analysis of *clear* is provided in (23).

- (23) *It is clear that p* is true just in case the maximal degree to which *p* is likely to be true is at least as great as  $\mathbf{d}(c)(\llbracket \text{clear} \rrbracket)$  [to be revised in (25)].

This analysis explains the connection between likelihood and clarity, and specifies the respect in which asserting clarity is similar to asserting the applicability of a vague predicate. However, it cannot be right. The problem is that in Stalnaker's model, propositions don't have probabilities. For any given possible world, either Mary is a doctor in that world or she isn't. This means that for any given world *c*, either the probability that Mary is a doctor is 1 or it is 0. Whatever the standard of clarity is, worlds in which the probability is 1 will survive update according to (23), and worlds in which the probability is 0 will not. But this is exactly the update effect of asserting *Mary is a doctor*: only those worlds in which Mary is a doctor will survive. Thus, the analysis in (23)

amounts to claiming that the meaning of *It is clear that Mary is a doctor* is identical to *Mary is a doctor*, which was shown above to be incorrect.

We believe the problem is solved by building on the observation that likelihood is a judgment made by some sentient creature who is contemplating  $p$ . If likelihood plays a role in assertions of clarity, we must figure out who is judging likelihood. An important clue comes from comparing a simple assertion of clarity to one in which the experiencer is overt:

- (24)a. It is clear that Mary is a doctor.
- b. It is clear to me that Mary is a doctor.
- c. (Surely) It is clear to you that Mary is a doctor.

The claim in (24a) is stronger than either of the claims in (24b) and (24c). In all cases, the speaker is committed to believing that Mary is a doctor, but (24b) allows the possibility that the hearer may not share that belief. With an implicit experiencer, as in (24a), there is a strong intuition that the experiencers of clarity must include at least the speaker and the listener (See Bhatt and Izvorski 1998 for arguments that (24a) has an implicit argument).

We approximate the meaning of (24a) as the conjunction of (24b) and (24c): if it is clear that Mary is a doctor, then it is clear to the discourse participants that Mary is a doctor. We refine our analysis of *clear* as in (25).

- (25) *It is clear to  $x$  that  $p$  is true in a world  $c$  just in case the maximal degree to which  $x$  judges that  $p$  is likely to be true is at least as great as  $\mathbf{d}(c)(\llbracket \text{clear} \rrbracket)$ .*

The revision considers judgments of likelihood at each world. That is, for any possible world  $c$ , how likely does the counterpart of  $x$  consider  $p$  to be? For instance, imagine that  $x$  is Gina Taranto, and  $c$  is a world in which the CIA is supremely devious and competent. They want Gina to think that Mary is a doctor, even though she is not, and they are so successful that Gina believes in  $c$  that Mary is a doctor. That is, the CIA conspiracy is effective, and this causes Gina to believe something that isn't true. In this situation, it is clear to Gina that Mary is a doctor, even though Mary isn't a doctor.

It is helpful to compare asserting clarity to asserting necessity, which is similar to asserting clarity, but which does not (directly) depend on belief. Compare (2) to a similar sentence with epistemic *must*:

- (26) Mary must be a doctor.

Both are guesses typically made on the basis of partial, indirect evidence. One key difference is that *must* does not implicate the existence of a so-called “judging experienter”, that is, a mind that judges what is abnormal versus what is expected. As a result, a speaker can assert (26) on the basis of private knowledge. In contrast, (2) requires that the hearer have access to all of the evidence necessary to come to the desired conclusion. Thus, any adequate analysis of *clear* must account for the public status of the evidence that provides the basis for the judgment.

A second, more subtle difference is that because *must* depends on what is normal or likely, there will always remain the possibility that something unlikely or abnormal happened and the conclusion doesn’t follow. This is why assertions of *must* are so often followed by requests for confirmation, as in (27).

(27) It *must* be a UFO or alien spacecraft, right?

Evidently, *must* does not commit the hearer to accept the designated proposition, or at least not very strongly, and the right of the hearer to doubt persists even if the hearer does not explicitly object. In contrast, once clarity has been asserted, failing to object immediately and firmly commits the hearer to accepting the truth of the relevant proposition. That is, if a speaker asserts *it is clear that Mary is a doctor*, and her hearer allows that assertion to go unchallenged, then the speaker is entitled to assume that the hearer believes that Mary is a doctor.

We propose that this difference between *must* and *clear* follows from the following fact: the truth conditions for *must* depend on examining worlds and their modal neighbors, and determining whether the proposition in question is true at those worlds. In other words, whether *must* holds depends on truth, while *clarity* depends directly on belief, and only indirectly on truth.

This claim is embodied in the analysis given in (25), provided we assume that  $\mathbf{d}(c)(\llbracket clear \rrbracket)$  returns the degree of likelihood required for someone to believe a proposition is true. That is, (25) recognizes that belief is a gradient attitude, and behaves just like any other vague predicate. For instance, the degree to which one believes Darwin is right may differ from the degree to which he believes an astrologer’s claim that retrograde motion of Mercury hinders communication.

In practical terms, this means that when a speaker asserts (28), the only worlds to survive update are those at which the speaker believes Mary is a doctor.

(28) It is clear to me that Mary is a doctor.

The surviving worlds will include every situation in which there is sufficient evidence to persuade the speaker that Mary is a doctor. Excluded worlds may include worlds in which the speaker knows Mary is on her way to a Halloween party, if this introduces enough uncertainty to reduce belief in her doctorhood to below the threshold specified by the delineation function for that world.

In particular, worlds may survive in which Mary is not a doctor, as long as the speaker believes Mary is a doctor in that world. In contrast with Beaver (2002), our claim is that clarity does not entail the proposition of which it is predicated. On our analysis the dialogue in (29) involves contradiction and repair, while the dialogue in (30) does not.

(29) A: Mary is a doctor.

B: Actually, Mary isn't a doctor. I asked her, and she revealed she's a CIA operative pretending to be a doctor.

(30) A: It is clear to me that Mary is a doctor.

B: Actually, Mary's not a doctor. I asked her, and she proved she's not.

In (29), we learn from B's contribution that A spoke falsely when A asserted that Mary was a doctor. In (30), however, B's statement does not contradict A: it remains true that it was clear to A that Mary was a doctor, so A spoke truly.

If asserting personal clarity does not entail that the proposition in question is true, how can we account for the fact that asserting simple clarity seems to entail the truth of the complement proposition? We suggest that the truth of the proposition is not in fact entailed. Rather, it is an illusion due to the implications that the assertion has for the state of the discourse. The chain of reasoning provided in (31) leads to the conclusion in (32).

(31) a. In the absence of an overt experiencer, the entities doing the believing default to both the speaker and the hearer.

b. The semantics of personal clarity guarantee that every world in the updated context will be a world in which the experiencer believes the truth of the proposition.

c. The result is that all of the discourse participants believe the truth of the proposition in every world in the updated context

(32) Therefore, asserting *it is clear that p* does not entail *p*, but guarantees the discourse participants are justified in behaving as if *p* is true.

This is another very peculiar situation. In terms of descriptive versus metalinguistic update, the update is entirely metalinguistic. To be more specific,

we have learned nothing new bearing on whether Mary is a doctor, since the only new information concerns the beliefs of the discourse participants. In particular, (33) lists some of the things that at least one of the discourse participants may not have known before asserting clarity that they would know after the assertion:

- (33) a. The speaker believes that Mary is a doctor.
  - b. The hearer believes that Mary is a doctor.
  - c. The speaker knows that the hearer believes that Mary is a doctor.
  - d. The hearer knows that the speaker believes that Mary is a doctor.
- [etc.]

(33a-b) are directly entailed by the proposed semantics of clarity. (33b) and (33c) (and the rest of the infinite regress) follow from the fact that the discourse participants assume that the other discourse participants agree to accept any assertion that goes unchallenged.

Importantly, the new information has nothing to do (at least not directly) with whether Mary is a doctor; its only effect involves the state of the discourse. Asserting clarity is about the judgment of the discourse participants, not about what is the case in the part of the world under discussion. Thus, asserting clarity synchronizes the common ground: it forces the speaker and the hearer to acknowledge that they are in a position to treat a proposition as if it were a fact.

## 5. Conclusions

Our analysis resolves our earlier paradoxes. Regarding Lemma 1 (if asserting clarity adds no new information about the situation under discussion, what use is it to assert it?), we claim that asserting clarity does add useful information about the state of the discourse— information about the attitude of the discourse participants towards the proposition in question.

Regarding Lemma 2 (if it is self-evident that Mary is a doctor, then isn't asserting clarity tantamount to suggesting that the hearer is an idiot?), we conclude that a speaker does not need to assume her hearer doesn't believe Mary is a doctor. It is sufficient for the speaker to assume the hearer may not know that the speaker also believes that Mary is a doctor.

Finally, regarding the reconstructed paradox (that *it is clear that p* is asserted only in situations in which it is in fact **not** clear that *p*), we conclude that asserting clarity does not require asserting perfect clarity: by recognizing the role of vagueness, we realize that asserting clarity means that the proposition is merely clear enough— in particular, clear enough to proceed as if it were true.

This understanding of the semantics of *clear* deepens the understanding of how context update works. In particular *clear* provides an example of a predicate whose meaning requires that the discourse model contain a model of itself. This is what we take to be Stalnaker's (1998) claim, though our implementation may go beyond what he explicitly advocated.

Furthermore, the case of clarity shows that there are expressions whose only update effect has to do with the state of the discourse, not the facts under discussion. This result is anticipated in recent work. Kyburg and Morreau (2002) show that some uses of vague expressions have the sole effect of negotiating vague standards. Additionally, Barker (2002) argues that there are constructions whose only discourse update effect is to negotiate vague standards. Asserting clarity is a much simpler and more direct case in which the only update effect is metalinguistic: asserting clarity provides information about the discourse and the discourse participants, and not about the facts under discussion.

## References

- Barker, Chris. 2002. "The Dynamics of Vagueness", *Linguistics and Philosophy*, **25**:1- 36.
- Beaver, David. 2002. *Presupposition*. Stanford: CSLI Publications.
- Bhatt, Rajesh, and Roumanya Izvorski. 1998. "Genericity, Implicit Arguments and Control". Proceedings of SCIL 7. MIT Working Papers in Linguistics.
- Fine, Kit. 1975. "Vagueness, Truth, and Logic", *Synthese*, **30**:265-300.
- Heim, Irene 1982. *The Semantics of Definite and Indefinite Noun Phrases*. PhD Dissertation, University of Massachusetts, Amherst.
- Kennedy, Chris. 1997. *The Syntax and Semantics of Gradability and Comparison*. PhD Dissertation, University of California, Santa Cruz.
- Kyburg, Alice and Michael Moreau. 2000. "Fitting Words: Vague Language in Context", *Linguistics and Philosophy* **23**.6:577-597.
- Lewis, David. 1970. "General Semantics", *Synthese*, **22**:18-67.
- van der Sandt, Rob. 1992. "Presupposition Projection as Anaphora Resolution", *Journal of Semantics* **9**:333-377.
- Stalnaker, Robert. 1979. "Assertion", in Cole, ed., *Pragmatics, Syntax and Semantics* **9**:315-332.
- Stalnaker, Robert. 1998. "On the Representation of Context", *Journal of Logic, Language, and Information*, **7**:3-19.
- Taranto, Gina. 2002a. "Discourse Adjectives", Paper presented at NASSLLI, Stanford, June 2002.
- Taranto, Gina. 2002b. "Experiencing Propositions". ms, UCSD, San Diego, California.
- Williamson, Timothy. 1994. *Vagueness*. London: Routledge.
- Williamson, Timothy. 1999. "On the structure of higher-order vagueness", *Mind* **108**:127-114.