The Paradox of Clarity: Defending the Missing Inference Theory

George Bronnikov*

SALT 18, UMass Amherst, March 20-23, 2008

1 The Problem, B&T's solution

In Barker and Taranto (2003), Taranto (2006), Barker (2007), construction *It is clear that p* is analyzed (as well as its variant *Clearly, p*).

As an initial approximation, the meaning of this construction seems to be that p is entailed by publicly available evidence. So, for example,

(1) It is clear that Abby is a doctor.

can be uttered when both the speaker and the hearer are looking at a picture of a woman wearing a lab coat and a stethoscope.

Barker and Taranto state the following problem: if the evidence presented to every participant of the conversation (part of the common ground) already entails p, there is no need in stating p. The common ground, viewed as a set of possible worlds, does not change after the assertion of clarity is made.

The solution proposed by Barker and Taranto involves the notion of a "linguistic side effect". Every sentence is assigned some truth conditions, and the dynamic effect of uttering it partly consists in narrowing the common ground by excluding those possible worlds that do not meet the truth conditions (this is the "main effect" of uttering the sentence). However, some changes to the common ground may not be related to the outside world, but to the state of the communication itself.[1] New discourse referents may be introduced. Standards may be set for vague predicates. For example,

(2) Bill is tall.

*University of Texas at Austin

[1]Barker and Taranto (2003), following Stalnaker, note that the conversation itself is part of the world. For this reason, they do not consider it necessary to add additional information to the common ground besides the set of possible worlds.

Barker (2007) treats the common ground as a set of pairs $\langle d, w\rangle$, where d is a state of the conversation (including standards for vague predicates), and w is a possible world. This difference plays no role in what follows.

may be uttered to provide information about Bill's height (that is, for its main effect; this would be descriptive use of *tall*), but it could also be uttered in order to specify what counts as tallness in the situation under discussion (metalinguistic use of *tall*).

Clarity assertions, according to Barker and Taranto, are always used exclusively for their side effects. Namely, they set standards for what evidence is considered sufficient for belief in their argument proposition p (according to Barker and Taranto 2003, Taranto 2006), or what evidence is considered appropriate justification for p (Barker 2007). That is, among the possible worlds constituting the common ground before the utterance, those worlds are excluded where the standards of belief/justification in the current conversation are set too high.

The theory proposed by Barker and Taranto has as its consequence that asserting the clarity of p does not in fact entail p.

2 Problems with B&T

There are, however, several problems with that theory. First, the prediction that *Clearly, p* does not entail p is not borne out. This can be easily seen by considering cases where p turns out to be false.

When clarity assertions is put in the present tense, B&T can predict the infelicity of statements like

(3) #It is clear that Abby is a doctor, but in fact she is not.

The clarity assertion in the first clause ensures that the speaker believes Abby to be a doctor. But in this case he cannot sincerely utter the second clause, otherwise he will fall victim to Moore's paradox.

However, as sson as we put the example in the past tense, those pragmatic factors are no longer at play. *Clear* examples (4) pattern with simple statements of a proposition (5), not with expressions of belief (6) or justifialbility assertions (7):

(4) a. #It was clear that Abby was a doctor, but in fact she was not.

b. It seemed clear that Abby was a doctor, but in fact she was not.

(5) a. #Abby was a doctor, but in fact she was not.

b. Abby seemed to be a doctor, but in fact she was not.

(6) a. We believed Abby to be a doctor, but in fact she was not.

b. It seemed to us that we believed Abby to be a doctor, but in fact she was not.

(7) a. It was justifiable to conclude that Abby is a doctor, but in fact she was not.

b. It seemed justifiable to conclude that Abby is a doctor, but in fact she was not.

In examples (4a) and (5a) we have a contradiction, which is absent in (6a) and (7a). In (4b) and (5b), the second clause denies correctness of the speaker's opinion expressed in the first clause. In (6b) and (7b) it does not.[2]

Secondly, in the Stalnakerian framework, once the standards of justification/belief are set, they can only get looser in the subsequent discourse (the worlds with tighter standards have already been eliminated). Consider, however, the following example.

(9) *A and B are sitting in an emergency room. A woman in a lab coat walks along the corridor.*
A: This is clearly a doctor.
A man walks by in the opposite direction. He wears a lab coat as well. He also has a stethoscope around his neck and carries a medical record under his arm.
A: Clearly, this is another doctor.

In Barker and Taranto's framework, the second utterance would be uninformative. The second person has more doctorlike features than the first one, and by the standards set in the first utterance, he should already be considered a doctor. The second utterance would thus be expected to be infelicitous. However it is perfectly normal.

Contrary to Barker and Taranto's claim, clarity assertions can be used in situations where there is no vagueness at all and the standards for belief/justification are completely determined. In particular, mathematical discourse:

(10) Take an integer n divisible by 9. Clearly, n is also divisible by 3.

Barker (2007) derives his analysis for cases where an experiencer is explicitly specified:

(11) It is clear to A that p

in the following way:[3]

> But if I assert *It is clear to me*, then I do not assume that you have encountered the same set of evidence. I am claiming, then, that if you knew what I knew – both in knowing the evidence in favor of the conclusion, and in failing to know any countervailing evidence that you might be aware of that I am not – then you would come to the same conclusion I have.

In this case, the clarity assertions lose their peculiar 'side-effect only' flavor: surely a statement

(12) It is clear to me that John is lying.

says something not just about the standards of evidence, but about the evidence itself that is available to me.

[2]Barker (2007) has an example similar to (4a):

(8) A: It is clear to me that Abby is a doctor.
B: No, actually, I asked her, and she said she's just an actor.
A: ($^{\#}$— Barker) Oh, I must have been mistaken.

My intuitions differ from his in this respect.

[3]He also states specifying an experiencer may mark personal preferences as to the required degree of justification.

3 Missing inference

My proposal is to take seriously the idea that the *clearly* construction marks the result of an inference. Namely,

(13) *It is clear to A from S that p* can be analyzed as: *A* has performed a valid inference which has *S* as premises and *p* as conclusion.

This is exactly what Barker (2007) calls the missing entailment theory (and dismisses). On my analysis, *It is clear that p* does entail *p*. By asserting clarity, the speaker takes full responsibility for the validity of his inference — even if the inference is defeasible. Thus, the behaviour of (4) is explained. In (9), the second utterance requires a separate (although similar) inference, so it is not superfluous.

Availability of clarity assertions for mathematical statements follows trivially from my account: these statements lose their special status; just like statements about the world, they are subject to inference operations.

If the *from S* part of the clarity assertion is omitted, then the evidence used as source for inference is left unspecified. In fact, when *A* is not the speaker, this inference may not be available to the speaker:

(14) I see that it is clear to you that John is lying; can you explain why?

If the *to A* part is left out, there should be an inference available to every participant of the conversation. Moreover, every participant should be able to make the same inference. This can be illustrated by the following example. Suppose John has read *Crime and Punishment*, Mary has read *The Brothers Karamazov*, and Peter *The Idiot*. When they gather to exchange their opinions, according to my intuition, it would not be appropriate for one of them to say:

(15) It is clear that Dostoevsky is a great novelist.

even though it is appropriate to utter

(16) It is clear to everyone here that Dostoevsky is a great novelist.

This kind of truth conditions requires that the speaker, in order to assert clarity of *p*, both be able to draw the inference himself and be able to attribute *the same* inference to the other participants. In order to attribute the inference to the other participants, he needs to know that they possess the premises *S* of the inference. This, of course, still comes short of the definition of the common ground (for example, the other participants may not know that the speaker knows that they know *S*), but it becomes rather hard to construct the tests, and when they are constructed, it is hard to elicit clear judgements on the appropriateness of using *Clearly, p* in such situations. So, for all practical purposes, my account predicts that the premises of the inference should be in the common ground when making clarity assertions without specifying the experiencer.

One way to capture the intuition that in a clarity assertion the speaker needs to have a specific inference in mind is to construct information states not just for

individual agents, but for groups as well. There is a discussion in von Fintel and Gillies (2007a) of ways to aggregate information states. For their purposes, however, an aggregated state is one where all the information possessed by a group is pooled (that would correspond to an intersection of possible world sets or to set union of information sets as sentences). In order to obtain the common ground, one would need to take into account only the information every participant in the conversation has (thus, set union of possible world sets or set intersection of representations).

Barker and Taranto's question 'why ever assert clarity?' receives a plausible explanation under this analysis: the speaker notifies the audience that the information they have (S) is sufficient to infer p. Each member of the audience is invited to build the inference for himself. The clarity statement can be used to build a greater confidence in the audience than simply stating p: upon deriving p, the hearer does not depend any longer on whether he trusts the speaker.

There are certain features noted by Barker and Taranto that any account of the *clearly* construction should be able to explain. Three of these features fall out immediately from my analysis. These are the inapplicability of clarity assertions to cases of direct evidence, information already explicitly stated in the conversation, and belief without proper justification (examples from Barker (2007)):

(17) #It is clear that Abby is wearing a stethoscope.

(18) A. Guess what? It turns out that Abby is a doctor!
B. #Now that you've told me this, it's clear that Abby is a doctor.

(19) #It is clear that God exists.

In all of these cases, there is no inference that allows us to assert p; therefore, the *clearly* construction is inappropriate.

Gradability of clarity assertions is a harder problem. There are several parameters by which inferences can be graded. Two are length of inference and likelihood of discovering it. However, as Barker's example:

(20) It is reasonably clear that Mars is barren of life.

shows, clarity is gradable with respect to the level of confidence that the inference provides to its conclusion. Most inferences in everyday life employ some amount of inductive and/or defeasible reasoning, so they don't guarantee the truth of their conclusion with absolute certainty. Moreover, people, with their limited reasoning capabilities, sometimes doubt whether the construction they have just built qualifies as an inference. Conclusions of inferences that are really bulletproof can be characterized as *absolutely clear*, *crystal clear*, and inferences that employ a lot of heuristics, generic reasoning and such can give rise to statements about propostitions that are *reasonably clear* or *relatively clear*. This analysis recovers much of the intuition behind Barker's theory. It also shows why mathematical inferences (even very long and complicated ones) hardly ever give rise to gradable clarity:

(21) ??It is reasonably clear that Fermat's Last theorem holds.

When sentences like (21) are used, this happens for the last of the reasons mentioned: the speaker does not have complete confidence that the proof he has in mind is correct.

4 Barker's objections

Barker's (2007) reasons for rejecting the "missing entailment" theory of the kind I am defending are twofold. First, clarity assertions are often made when the proposition in question is not in fact entailed by the evidence:

(22) It is clear that Abby is a doctor

is said when she might in fact be an actress or dressed for a Halloween party. All we need to say is that inference, the existence of which is stated by the clarity assertion, may be defeasible: it can involve generic statements or other types of default rules. There will be no purely logical entailment in such cases.

Second, for some examples the missing entailment theory seems to predict wrong results:

(23) A. John is a bachelor.
B. #Clearly, then, he is not married.

(24) A. John ate a sandwich and drank a glass of beer.
B. #So it is clear that John ate a sandwich.

We can note that the inferences involved in these examples are extremely simple: subtyping in (23) and conjunction simplification in (24). So the missing entailment theory can be saved if we specify that the entailment in question should be substantial enough — not limited to certain easy types of inference.

5 A formal theory

So far, all our arguments were informal. To make them more precise, we need a framework where one can take inference into account. Unfortunately, the commonly accepted semantics based on possible worlds and accessibility relations is not such a theory, since it paints a picture of logically omniscient agents. Instead, we can resort to a picture where the cognitive state of an agent is represented by a set of sentences.

A state in our model consists of a world, a finite set of sentences representing the speaker's beliefs[4], and a finite set of sentences representing the hearer's beliefs. For any formula ϕ of our logic language and any agent a, $\mathbf{B}_a\phi$ will be a formula which is true just in case ϕ is in a's belief set.

We can use Duc's (2001) idea and employ a variant of dynamic logic to represent inferences. An application of a derivation rule R by an agent a constitutes an action; this action converts a state into one where the conclusion of the rule is added to a's cognitive state.

Duc is only interested in computational complexity of derivations; thus, his logic only has F_a as an action — that is, application of *any* rule. We need to talk about

[4]Beliefs of an agent are not limited to sentences in this set. We need to account for tacit beliefs such as *There are less than 50000 people in this room.*

derivations differing in how obvious they are; presumably, some rules are harder to apply than the others. Therefore, we consider every rule as a separate type of action.[5]

In order to provide an analysis of clarity assertions, we need to distinguish between trivial derivations (which are too simple to trigger clarity assertions), easy derivations (which do trigger asserions of clarity) and hard derivations (where the existence of such a derivation does not count as clarity). A simplifying assumption here will be that these types of derivation only differ in the types of rules employed.[6] Thus for an agent a we will have trivial rules $A_{1a} \dots A_{ka}$, easy rules $B_{1a} \dots B_{ma}$, and hard rules $C_{1a} \dots C_{na}$. A trivial derivation performed by a will be an action $Triv_a = (A_{1a} \cup \dots \cup A_{ka})^*$ (an action composed of any number of elementary actions, where each elementary action is an instance of a trivial rule applied by a). An easy inference by a is the action $Easy_a = (A_{1a} \cup \dots \cup A_{ka} \cup B_{1a} \cup \dots \cup B_{ma})^*$, that is, it can employ both trivial and easy rules any number of times.

In this case, we can express *It is clear to a that* ϕ as

$$\langle Easy_a \rangle \mathbf{B}_a \phi \wedge \neg \langle Triv_a \rangle \mathbf{B}_a \phi$$

For example, assume that conjunction simplification (CS) is a trivial rule, and universal exploitation (UE) and modus ponens (MP) are easy rules.

Suppose an agent a is in the following information state:

$$S_1 = \left\{ \begin{array}{c} N \bmod 9 = 0, \\ \forall x(x \bmod 9 = 0 \rightarrow x \bmod 3 = 0) \end{array} \right\}$$

In this state, it will be true that $\mathbf{B}_a(N \bmod 9 = 0)$ Since by applying rules UE and MP, a can achieve the state

$$S_2 = \left\{ \begin{array}{c} N \bmod 9 = 0, \\ \forall x(x \bmod 9 = 0 \rightarrow x \bmod 3 = 0), \\ N \bmod 9 = 0 \rightarrow N \bmod 3 = 0, \\ N \bmod 3 = 0 \end{array} \right\}$$

the following formulas will be true in S_1:

$$\langle \mathrm{UE}_a; \mathrm{MP}_a \rangle \mathbf{B}_a(N \bmod 3 = 0)$$
$$\langle (\mathrm{UE}_a \cup \mathrm{MP}_a)^* \rangle \mathbf{B}_a(N \bmod 3 = 0)$$
$$\langle Easy_a \rangle \mathbf{B}_a(N \bmod 3 = 0)$$

Since $\langle Triv_a \rangle \mathbf{B}_a(N \bmod 3 = 0)$ is false in this situation, (10) is true, according to our definition.

[5] A model like this is limited to one particular derivation system. However, systems may be intertranslatable. Derivation patterns using one system of rules may be expressible as derivation patterns using another system.

If we use natural deduction as our system of rules, we will need to represent subproofs as action functions (functions taking an action and producing another). This, in turn, necessitates the use of Context Free Dynamic Logic (Harel 1977).

[6] In principle, the number of steps and their order could also play a role, and such conditions are also expressible via the language of dynamic logic.

6 *Clearly* vs. epistemic *must*

In vonFintel and Gillies (2007b), an argument similar to mine is made with respect to the epistemic *must*, and a similar solution is proposed:

> Epistemic modals signal that their prejacent is not directly settled by the salient kernel (where 'kernel' is a non-logically closed set of sentences – *G. B*).

However, *clearly* and *must* are not interchangeable.

- In the *clearly* construction, the existence of an appropriate inference is part of the assertion. Unlike *must*, *clearly* can take narrow scope with respect to operators like negation and tense.

  (25) It is not clear to me that Abby is a doctor, but she might be.

  (26) It was clear to me yesterday already that Abby is a doctor.

- *Must* does not have to be based on public evidence, even when the relevant group is not specified explicitly.

- Certain types of inference can be marked by *must*, but not by *clearly*:

  (27) *John left two hours ago*
  a. He must be home by now.
  b. ?Clearly, he is home by now.

- One can use *clearly* (but not *It's clear that*) to signal an inference whose conclusion is already known to the speaker.

  (28) Mary has been out of town for three days. She has not phoned. Clearly, I'm worried/#I must be worried.

7 Final remarks

Using explicit representations of the world in the mind of an agent is a *conservative* extension of the standard semantics based on possible worlds, since every such representation uniquely determines a set of possible worlds where all its statements are true.[7]

Representation of sentences in an internal language, manipulated by inference, is a philosophically plausible idea (Fodor 1975 is perhaps the most famous exposition).

It could be, however, that semantics of natural language never refers to those representations, just to possible world structures induced by them. The existence of clarity assertions shows that this is not the case.

[7]Provided that the question of reference for elementary symbols of the internal language is resolved.

Other constructions where a theory that deals with inference explicitly may prove useful include indirect speech, belief ascriptions and evidentials (both hearsay and inferential).

In the case of indirect speech and hearsay evidentials, the words of the primary speaker can undergo certain operations that can be called inference, such as dropping of arguments or adverbials:

(29) John: I sold my car to Bill yesterday. Peter: John says that he sold his car.

In the case of belief ascriptions, certain limited inferences have to be allowed to account for the so called tacit beliefs (see e. g. Konolige 1986):

(30) John believes that there are less than 50000 people in this room.

In the case of inferential evidentials, the type of inference they allow may be limited (inference to process from its results; generic inference).

In all these cases, the language of patterns representing shapes of deductions may prove useful.

8 References

Barker, Chris. 2007. Clarity and the grammar of skepticism. `http://semanticsarchive.net/Archive/zExYWRkY/barker-clarity.pdf`

Barker, Chris, and Gina Taranto. The paradox of asserting clarity. In Paivi Koskinen (ed.), *Proceedings of the Western Conference on Linguistics (WECOL) 2002*, Volume 14, Department of Linguistics, California State University, Fresno. 10–21.

Duc, Ho Ngoc. 2001. *Resource-Bounded Reasoning About Knowledge*. Ph. D. thesis, Univ. of Leipzig.

Fodor, Jerry. 1975. *The Language of Thought*, Harvard University Press.

Harel, David. 1979. *First-Order Dynamic logic*. Berlin: Springer-Verlag.

Konolige, Kurt. 1986. *A Deduction Model of Belief*. San Francisco: Morgan Kaufmann.

Taranto, Gina. 2006. *Discourse Adjectives*. NY:Routledge.

von Fintel, Kai, and Anthony S. Gillies. 2007a. 'Might' Made Right. `http://mit.edu/fintel/mmr.pdf`.

von Fintel, Kai, and Anthony S. Gillies. 2007b. *Must*...Stay...Strong! `http://mit.edu/fintel/mss-handout.pdf`