

The Paradox of Clarity: Defending the Missing Inference Theory

George Bronnikov
University of Texas at Austin

1. The Problem, B&T's solution

In Barker and Taranto (2003), Taranto (2006), Barker (2008), construction *It is clear that p* is analyzed (as well as its variant *Clearly, p*).

As an initial approximation, the construction seems to mean that *p* is entailed by the evidence available to some relevant group, which typically includes all the participants of the conversation. 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.¹

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.² New discourse referents may be introduced. Standards may be set for vague predicates. For example,

- (2) Bill is tall.

I would like to thank Nicholas Asher, Derek Ball, John Bengson, Gennaro Chierchia, Kai von Fintel, Bryan Pickel, Malte Willer, and the anonymous reviewers for helpful comments on my work. I am especially grateful to Chris Barker for getting me interested in the subject.

¹As (Barker 2008) notes, there are cases where the set of relevant participants does *not* contain both the speaker and the audience. In these circumstances the paradox does not arise. One type of such cases are assertions of personal clarity.

²Barker 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 information to the common ground apart from the set of possible worlds.

Barker (2008) 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 is the formalism I use in discussing B&T's theory. The difference plays no role in what follows, however.

may be uttered to provide information about Bill's height (that is, for its main effect; this would be a descriptive use of *tall*), but it could also be uttered in order to specify what counts as tallness in the situation under discussion (a 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 2008). 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 it as its consequence that asserting the clarity of p does not in fact entail p .

2. Problems with B&T

2.1. Factivity

There are, however, some problems with this 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.

Considering cases where clarity assertions stand 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 she cannot sincerely utter the second clause, on pain of falling victim to Moore's paradox.

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

Barker claims that *Clearly, p* patterns with belief assertions (non-factive) rather than knowledge assertions (factive), since they can be combined with *might not p* claims without contradiction. My intuitions differ from his on this point. Substituting an actual sentence for *p* in his examples to ease judgment, we get

- (8) a. We know that Abby is a doctor, although she might not be.
 b. We believe that Abby is a doctor, although she might not be.
 c. It is clear that Abby is a doctor, although she might not be.

It seems to me that the only way to avoid inconsistency in (8c) is by making a pause after the first clause:

- (9) It is clear that Abby is a doctor. . . wait a minute, she might not be one, she might be an actress.

One can repair hasty knowledge claims in a similar way. So clarity assertions pattern with factive statements after all.

2.2. Repeated clarity assertions

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

- (10) *A and B are sitting in an emergency room. A woman in a lab coat (X) walks along the corridor.*

A: This is clearly a doctor.

A man (Y) 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.

Suppose we have four possible worlds:

- w_1 X and Y are both doctors.
 w_2 X is a doctor. Y is not.
 w_3 Y is a doctor. X is not.
 w_4 Neither is a doctor.

We also have three possible degrees of skepticism (these are part of the state of conversation; we are not interested in the other parts):

- d_1 Wearing a lab coat is sufficient to be judged a doctor.
 d_2 Wearing a lab coat is not sufficient, but together with a stethoscope and a medical record it does satisfy our doubts.
 d_3 Nothing, even the medical record, is convincing enough.

Note that, since Y has more doctorlike features, there is no refinement of the vague standard for justification that would make X count as a doctor, but not Y.

At the start of the conversation, every world/standards combination is possible:

$$S_1 = \left\{ \begin{array}{l} \langle d_1, w_1 \rangle, \langle d_1, w_2 \rangle, \langle d_1, w_3 \rangle, \langle d_1, w_4 \rangle, \\ \langle d_2, w_1 \rangle, \langle d_2, w_2 \rangle, \langle d_2, w_3 \rangle, \langle d_2, w_4 \rangle \\ \langle d_3, w_1 \rangle, \langle d_3, w_2 \rangle, \langle d_3, w_3 \rangle, \langle d_3, w_4 \rangle \end{array} \right\}$$

After the first utterance, those world-standard pairs that don't allow lab coat to count as enough evidence for doctorhood are eliminated (note that the clarity assertion does not tell us anything about the world itself):

$$S_2 = \{ \langle d_1, w_1 \rangle, \langle d_1, w_2 \rangle, \langle d_1, w_3 \rangle, \langle d_1, w_4 \rangle, \}$$

The second clarity assertion could serve to eliminate d_3 out of the set of possible standards, but these world-standard pairs are already eliminated by the time it is uttered. Thus, in Barker and Taranto's framework, the assertion would be uninformative and therefore infelicitous. However it is perfectly normal.

2.3. No vagueness

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:

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

To accommodate these cases, B&T would have to argue either that there are vague standards of belief/justification involved after all (in particular, that there are context elements where n is divisible by 9, but somehow not by 3), or that this kind of use is special and needs separate treatment. If they choose the latter option, an explanation would be in order, first, why the theory for mathematical (and similar) uses of *clearly* does not apply to the more mundane situations, and second, why the polysemy of the *Clearly, p* construction is the same across a wide variety of languages.

3. Missing inference

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

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

This is exactly what Barker (2008) 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 soundness of her inference — even if the

inference is defeasible. Thus, the behaviour of (4) is explained. In (10), the second utterance requires a separate (although similar) inference, so it is not superfluous.

Availability of clarity assertions for mathematical statements follows trivially on 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:

(13) 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:

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

even though it is appropriate to utter

(15) 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 herself and be able to attribute *the same* inference to the other participants. In order to attribute the inference to the other participants, she 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 judgments 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 (2007) 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 states as sets of 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).

The requirement for the inference justifying the clarity assertion to be sound can be underscored by Gettier cases: suppose Abby is in fact a doctor, but she is dressed in a lab coat for a Halloween party, with a toy stethoscope around her neck.

Under this scenario, the inference “lab coat means doctorhood” is not sound, and (1) is false.

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 (2008)):

- (16) #It is clear that Abby is wearing a stethoscope.
- (17) 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.
- (18) #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.

4. Barker’s objections

Reasons given in Barker (2008) for rejecting the “missing entailment” theory of the kind I am defending are the following. First, clarity assertions are often made when the proposition in question is not in fact entailed by the evidence:

- (19) 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 as premises or use other types of default rules. There will be no strict entailment in such cases.

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

- (20) A. John is a bachelor.
B. #Clearly, then, he is not married.
- (21) 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 (20) and conjunction simplification in (21). 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.

The requirement that the inference be nontrivial may stem from the fact that people are reluctant to recognize certain simple inference steps as such.³ After all,

(22) John ate a sandwich and drank a glass of beer. [?]Therefore, he ate a sandwich.

does not sound too natural either. If we do recognize absence of trivial inference as part of the meaning of *clearly*, it has to be a presupposition, as demonstrated by a negation test:

(23) A. John is a bachelor.
B. #It isn't clear that he is not married.

(only allowable if B is disputing A's claim).

The class of 'trivial' inferences may, for all I see, coincide with the class of inferences involved in ascribing beliefs to other persons:

(24) a. Bill believes that John is a bachelor. Therefore, he believes that John is unmarried.
b. Bill believes that n is divisible by 9. #Therefore, he believes that n is divisible by 3.

In (24a), we have a valid inference. It is impossible for a competent speaker to believe someone to be a bachelor without believing him to be unmarried. The inference in (24b), on the other hand, is invalid, since humans are not logically omniscient.

Gradability, which Barker uses as another argument against the missing entailment theory, is discussed later, in section 8.

5. A formal theory

So far, all our arguments have been informal. To make them more precise, we need a framework where one can take inference into account. Inference consists in manipulating representations, so I choose to introduce representations explicitly into our model.⁴ The cognitive state of an agent a (in each possible world) will be captured by a set S_a of formulas — her explicit beliefs.⁵ These formulas are supposed to be translations from some internal "language of thought", free of indexicals. 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. To capture group knowledge, we may introduce sum agents: $\mathbf{B}_{a+b}\phi$ will be true iff ϕ is in the belief sets of both a and b .

³This explanation was suggested by Gennaro Chierchia (p. c.).

⁴See section 6 for a criticism of an approach which tries to deal with inference using a possible worlds toolbox.

⁵Beliefs 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.*

Epistemic accessibility relation between worlds (for an agent a) can be recovered by taking the intersection of propositions denoted by the sentences in a 's belief set.

Admittedly, “sentence” is not a good term for the internal representation of a belief. These “sentences” are more like the structured propositions of (Lewis 1970) or thoughts in (Crimmins 1992). I stick with “sentence” as a term since “proposition” is reserved for sets of possible worlds, and talking about inferences on ‘thoughts’ seem too unconventional.⁶

We can use an idea stated in Duc (2001) and employ a variant of dynamic logic to represent inferences.⁷ 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.⁸

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 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.⁹

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 assertions 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. (In principle, the number of steps and their order could also play a role, and such conditions are also expressible in the language of dynamic logic.) 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.

⁶Thoughts in the sense of Crimmins are different from structured propositions: their basic building blocks are notions and concepts, not word denotations. This setup makes it possible to deal with failures of reference for elementary symbols of the internal “language”.

⁷Another related system is Logic of Proofs (see, for example, Artemov 2001).

⁸In Dynamic Logic, $\alpha; \beta$ denotes an action consisting of sequentially performing α , then β ; $\alpha \cup \beta$ is non-deterministic choice between α and β ; and α^* is action α repeated 0 or more times. $\langle \alpha \rangle \phi$ is a formula saying that ϕ *may* hold after α is performed. Since we only distinguish our elementary actions up to the rule being applied (i. e. it may be applied to the wrong premises), we have no use for the stronger $[\alpha] \phi$ form.

⁹A 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 1979).

In this case, we can say that

$$\llbracket \text{It is clear to } a \text{ that } \phi \rrbracket = \langle \text{Easy}_a \rangle \mathbf{B}_a \phi$$

with the presupposition

$$\neg \langle \text{Triv}_a \rangle \mathbf{B}_a \phi$$

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:

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

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

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

the following formulas will be true in S_1 :

$$\begin{array}{l} \langle \text{UE}_a; \text{MP}_a \rangle \mathbf{B}_a(N \bmod 3 = 0) \\ \langle (\text{UE}_a \cup \text{MP}_a)^* \rangle \mathbf{B}_a(N \bmod 3 = 0) \\ \langle \text{Easy}_a \rangle \mathbf{B}_a(N \bmod 3 = 0) \end{array}$$

Since $\langle \text{Triv}_a \rangle \mathbf{B}_a(N \bmod 3 = 0)$ is false in this situation (and thus, the presupposition for a clarity assertion is satisfied), (11) is true, according to our definition.

6. *Clearly* vs. *epistemic must*

In von Fintel and Gillies (2008), 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 set of propositions — a structure that does not have the closure property — 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.

(This is a property that many epistemic modals have, but by no means all of them. For example, the English *have to* can be embedded under tense and negation operators.)

- *Must* does not have to be based on public evidence, even when the relevant group is not specified explicitly. In fact, there is no way to specify the relevant group in *must*.
- In certain situations, an inference can be marked by *must*, but not by *clearly*:

(27) *John left two hours ago. Every participant in the conversation knows that.*

- a. He must be home by now.
- b. ?Clearly, he is home by now.

Note, however, that once the premises of the inference are stated explicitly, clarity assertion becomes better:

(28) John left two hours ago. It takes only half an hour to get home from here. Clearly, he is home by now.

I can see two ways to explain this behaviour:

1. Clarity assertions require the premises of the inference to be actively entertained by the participants of the conversation. In (27), the information required to deduce that John must have arrived by now sits in the background of the interlocutors' minds. Once it is foregrounded in (28), one can use the *Clearly, p* construction.

The assumption that inferences based solely on background information do not give rise to clarity seems to be refuted, though, by the following example:

(29) The economy is clearly in recession.

can be uttered "out of the blue", without any preparatory foregrounding statements.

2. In (27), the "derivation" that leads to the conclusion involves operations that do not look very much like traditional inference steps: something like constructing mental scales and measuring distances on those scales. Presumably, such activity does not count as "easy" for the purposes of clarity assertion. On the other hand, in (28), the premises are stated linguistically, and the inference consists in a couple of standard rule applications.

- One can use *clearly* to signal an inference whose conclusion is already known to the speaker.

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

What matters in this example is that an inference exists from public evidence that leads to the conclusion stated in the prejacent. Even though the speaker has more direct means of knowing the prejacent, the use of *clearly* is sanctioned.

As for the solution proposed by von Stechow and Gillies, it involves (as has been noted already) contexts as sets of propositions. Such a set induces a partition on possible worlds. Proposition p is not settled by the context if there are possible worlds belonging to the same class in the partition which don't agree on p . Such a construction does not distinguish between equivalent propositions, so it is easy to construct counterexamples to the theory using standard philosophical test cases:

- (31) a. This animal has a heart.
 b. So, clearly, it has a kidney as well/So it must have a kidney as well.
- (32) a. Triangle ABC is equilateral.
 b. Clearly, it is equiangular/So it must be equiangular.

Assuming that creatures with a heart are necessarily all and only creatures with a kidney, proposition *this animal has a heart* is exactly the same as proposition *this animal has a kidney*. Upon uttering (31a), this proposition is settled in the context of the conversation. (31b) (in either of its variants) should be abnormal. In fact, it is fine.

7. Special case: Sherlock Holmes

There is one special use case of clarity assertions, which can be demonstrated by an example, suggested by Derek Ball (p. c.). Imagine Sherlock Holmes investigating a case together with Dr. Watson and Inspector Lestrade. After all the evidence has been collected (and known by all participants in the conversation), Holmes points out the murderer:

- (33) It is clear that the butler did it. For the maid was out on leave on the day of the murder, the gardener is deaf and would not hear the doorbell. . . etc. etc.

What distinguishes this kind of usage from the standard one is that the inference that leads to the conclusion can be arbitrarily complex. After making the clarity assertion, the speaker immediately presents the inference.

Perhaps the simplicity of inference is just a pragmatic requirement in the standard clarity assertion cases. When the inference is not presented immediately by the speaker, the clarity assertion seems pointless as long as the audience is not

able to recover it: the assertion does not increase the hearer's confidence in the proposition stated.

8. Gradability

As emphasized by Barker (2008), clarity is gradable: we have expressions like *crystal clear*, *somewhat clearer* etc. The theory presented in this article does not allow us to capture this property of clarity assertions. I have to resort to an informal description as to where the sources of gradability might be located.

There are several parameters by which inferences can be graded. Two are the length of inference and the likelihood of discovering it. As Barker's example:

(34) 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 derivation they have just built qualifies as a valid (much less sound) 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 propositions 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:

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

When sentences like (35) are used, this happens for the last of the reasons mentioned: the speaker does not have complete confidence that the proof she has in mind is correct.

The fact that clarity is gradable shows that we are unlikely to discover one day exactly what pattern of inference can give rise to clarity assertions: this is context dependent and vague. This vagueness can lead to side effects of the sort described by Barker: a clarity assertion may serve to establish which inferences count as easy. However, just like in cases with *tall*, such side effects do not exhaust the meaning of the construction.

9. 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.

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.

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:

- (36) 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)):

- (37) 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 is likely to be useful.

References

- Artemov, Sergei N.: 2001, 'Explicit provability and constructive semantics', *Bulletin of Symbolic Logic* 7, 1–36.
- Barker, Chris: 2008, 'Clarity and the grammar of skepticism'. In press at *Mind and Language*.
- Barker, Chris and Gina Taranto: 2003, 'The paradox of asserting clarity', in P. Koskinen (ed.), *Proceedings of the Western Conference in Linguistics (WECOL) 2002*, Vol. 14, 10–21. Department of Linguistics, California State University, Fresno, CA.
- Crimmins, Mark: 1992, *Talk About Beliefs*. MIT Press.
- Duc, Hon Ngoc: 2001, *Resource-Bounded Reasoning About Knowledge*, Doctoral Dissertation, University of Leipzig.
- von Fintel, Kai and Anthony S. Gillies: 2007, 'Might Made Right'. ms., MIT and University of Michigan. Available at <http://mit.edu/fintel/mmr.pdf>.
- von Fintel, Kai and Anthony S. Gillies: 2008, 'Must...Stay...Strong!'. ms., MIT and University of Michigan. Available at <http://mit.edu/fintel/fintel-gillies-2008-mss.pdf>.
- Fodor, Jerry: 1975, *The Language of Thought*. Harvard University Press.

- Harel, David: 1979, *First-Order Dynamic Logic*. Springer-Verlag, Berlin.
- Konolige, Kurt: 1986, *A Deduction Model of Belief*. Morgan Kaufmann, San Francisco.
- Lewis, David: 1970, 'General Semantics', *Synthese* **22**, 18–67.
- Taranto, Gina: 2006, *Discourse Adjectives*. Routledge, NY.