The reference argument of epistemic *must*

Matthew Stone*
Computer and Information Science
University of Pennsylvania
200 South 33rd, Philadelphia PA 19104-6839
phone: (215) 898-9517
e-mail: matthew@linc.cis.upenn.edu

August 1994

Abstract

Epistemic *must* is used to present a conclusion. In this paper, I explore the hypothesis that this should be modeled computationally using the notion of ARGUMENT presented by Simari and Loui [16]. An utterance of **must** p in conversational context κ is interpreted as asserting that the argument $\langle A, p\rangle$ is justified in κ. The parameter A provides a set of reasoning rules which, along with factual premises from which they derive p, must be salient in κ for the utterance to be felicitous.

Simari and Loui's formulation describes a relationship of defeat between arguments. Thus, in this account as in previous ones, the conclusions presented by epistemic *must* may be defeasible. This proposal improves on previous accounts in three key respects. First, the criterion that the argument be justified ensures that the speaker believes p when uttering **must** p. Second, the requirement that the speaker intend the hearer to recover the argument helps to explain the distribution and of *must* in discourse and the accommodation sometimes involved in understanding uses of *must*. Third, the link between the claim made by *must* and a specific argument correctly predicts the variation in apparent force of the modal in different contexts: it varies according to the strength of the argument and the speaker's intentions in providing the argument.

Because this interpretation for *must* incorporates restrictions based on salience into a framework designed to be relatively tractable, it may be uniquely suited for implementation.

1. Introduction

Consider the modality of EPISTEMIC NECESSITY introduced by *must* in (1).

(1) That must be Harry at the door. (= [20] 5.27)

In describing the meaning that this epistemic *must* contributes to a sentence, a principle along the lines of (2) is often proposed.

(2) **must** p, in contrast to p, PRESENTS A CONCLUSION.

For example, Ehrman [3] offers the explication that for me to utter **must** p for some predication p means that "the predication is required by my view of the probable CONSEQUENCES [emphasis added] of all the relevant factors (68)." Meanwhile, Woisetschlaeger [20] has: "It is a fact about English that a sentence like (27) [(1)] reports an inference based on PRESENTLY AVAILABLE EVIDENCE (115)." Coates [1] writes:

*Thanks to Robin Clark and Mark Steedman for discussion. This research was supported by an NSF graduate fellowship.

> In its most normal usage, Epistemic *must* conveys the speaker's confidence in the truth of what he is saying, based on a deduction from facts known to him (which may or may not be specified). (41)

And to Coates, Palmer [14] provides a refinement:

> Moreover, it is the notion of deduction or inference from known facts that is the essential feature of *must*, not just the strength of commitment by the speaker. For *must* does not have the same kind of meaning as the adverbs *certainly*, *definitely*, etc., which are, indeed, indications simply of the speaker's confidence or commitment. (59)

This paper offers a novel formalization of these ideas. Its contribution is in carefully describing the kind of presentation involved in using epistemic *must*. This description is based on the idea that *must* presents a conclusion by referring to the argument that justifies the conclusion. In other words, **must** p should be paraphrased as in (3).

(3) Some particular collection of facts A, salient in the common ground, provide (or have provided) a decisive reason to adopt the belief that p.

Simari and Loui's ARGUMENT SYSTEMS [16] provide the formal machinery with which to implement the notion of "providing a decisive reason". Argument systems are a framework for defeasible reasoning, so this formalization can describe the complex patterns of inference which have been captured in previous formal accounts of epistemic *must*. Because of its emphasis on presentation, however, this account improves previous ones by correctly explaining facts about the accommodation and implicature associated with uses of epistemic *must* in discourse.

2. A Different Approach

In formal semantics, the meaning of **must** p is often elucidated using a paraphrase like (4) (adapted from Kratzer [11] for example).

(4) In view of what is known, it is necessary that p.

There is a problem with using this paraphrase to predict the distribution of *must*. Following Grice's Maxim of Quality [7], a speaker may assert p only if he has evidence for it, that is, only if he knows it. But, in usual logics of knowledge, one knows everything that follows from what one knows. Thus whenever (4) can be asserted, p can be asserted too: there seems no reason to use *must* at all.

Semantic analyses based on (4), then, need a mechanism that semantically distinguishes conclusions that can be presented directly from those that can only be presented using *must*. Two such analyses are widely known: the data semantics of Landman [12] and Veltman [18], and possible-worlds semantics as enriched by Kratzer [10, 11].

For data semantics, the key is a model-theoretic realization of information that distinguishes those facts that are explicit in a speaker's information-state from those that are merely implicit. A collection of consistent pieces of information explicitly supports some fact if the fact is literally a member of the collection; in contrast, it implicitly supports some fact if the fact will remain consistent with the information no matter how the information may be consistently extended. The claim of data semantics, then, is that p is true in information-state s just in case p is explicit in s, while **must** p is true in s just in case p is implicit in s.

For possible-worlds semantics, the key is describing modal inference in terms of two structures: a MODAL BASE and an ORDERING SOURCE. The modal base associates a set of accessible alternatives with each world; the epistemic modal base appropriate for (4) gives for world w the set of worlds where what is known in w is true. The ordering source provides a way of comparing these worlds, perhaps for (4) on the basis how normal for w the events which take place in them are. **must** p is then true in w provided the p is true in the worlds in the modal base that are maximal by the ordering source. On this approach, it is possible for **must** p to be true when p is not because the ordering source allows more information than just what is known, such as what normally happens, to be taken into account in concluding p. Meanwhile, p is true only if it is a logical consequence of what one knows.

3. No Special Conclusions

These semantic proposals invoke the idea that people explicitly signal the different logical status of different conclusions in conversation. This claim should be met with skepticism: conventional wisdom, as articulated by Harman [9], is this:

> It would seem that one can RECOGNIZE a logical implication or logical inconsistency only if one has the relevant concept of logical implication or logical inconsistency. But it would seem that few people have such concepts, at least if this involves distinguishing logical implication and inconsistency from other sorts of implication and inconsistency. (17)

Indeed, it is quite difficult to characterize modal and indicative assertions in naturally-ocurring data on the basis of the kind of inference involved. Some examples from the Brown corpus [5] illustrate the problem.

All agree that indicative assertions are appropriate in presenting strictly logical inferences. (5), in which the second clause offers a mathematical consequence of the first, exemplifies this:

(5) The travel time of sound in tissue is about 1500 meters per second; thus it takes about 16 msec. to traverse 25 mm. of tissue.

However, there is no class of conclusions so direct or incontrovertible that only indicative assertions may present them. In fact, one of the most common uses of epistemic *must* comes from mathematical proofs, where every step of reasoning meets the strictest criteria of validity:

(6) Suppose [Af][1] is defined in the sub-interval [Af]. Now [Af] and [Af] must both be tangent points on the T component in the f-plane; otherwise by Lemma 1 the component would extend beyond these points.

Nor is there any class of conclusions so indirect or tenuous that they can only be presented using epistemic *must*. In presenting conclusions generally, *must* can be used to infer the supposed consequences of a situation, as in (7).

(7) People must think the curse is on me, seeing you fresh as an apple and me old and gray.

Alternatively, *must* may suggest inference to the best explanation of a result situation, as in (8).

(8) The handsome bird was solitary; its mate must be at home, silently guarding their nest.

Yet the very same possibilities can be expressed with the indicative, under the right conditions. The association in (9) makes an inferential connection between a state and its likely consequence.

(9) ...I am no longer young and [am] therefore incapable of being youthful of mind.

Meanwhile, (10) describes an explanation for an unexplained silence.

(10) Something else distracted him, yet there was no sound, only tomblike silence. Then he knew it was not sound, but lack of it. The air conditioner was no longer running.

Speakers express a variety of conclusions, some logical, some defeasible. Sometimes they express them using epistemic *must*, but the data suggests that this use does not reflect the different logical status of conclusions.

4. Implicature and Accommodation

I propose instead that speakers choose to use epistemic *must* when it serves their communicative intentions to make the dependence of claims on evidence particularly clear, not when their claims depend on evidence in a special way. Striking evidence for this comes from patterns of accommodation and implicature associated with expistemic *must*.

For example, consider the contrast between use of the word *should* and parallel use of the word *must*, in examples (11) and (12). (Woisetschlaeger discusses a similar contrast [20].)

(11) John asked, "Where is the sugar?" and Bill answered, "It should be in the cabinet over the fridge."

[1] Formulas have been deleted in the Brown corpus and replaced with "Af".

(12) Ann asked, "Where is the sugar?" and Mary answered, "It must be in the cabinet over the fridge."

Imagine stumbling across either of these sentences as part of a story. Understanding them involves explaining how come the first speaker asks this question, and how come the second arrives at this answer. When the discourse is presented out of context, this interpretation triggers the accommodation of several assumptions. For example, you can infer from John's and Ann's question that they are carrying out some activity for which they need sugar and that they find themselves in circumstances where they expect their addressee to know where the sugar is.

The assumptions accommodated for Bill's reply are quite different from those accommodated for Ann's. Bill's reply seems to invoke a modality of epistemic necessity along the lines suggested by (4): Bill answers that, if everything is normal, in view of his knowledge, the sugar is in the cabinet over the fridge. To understand Bill's utterance, the reader need only accommodate some reason for Bill to know where the sugar normally is. You might suppose that John is at Bill's house making cookies (this explains the question) and that Bill normally puts the sugar over the fridge, but sometimes accidentally exchanges it with flour which goes in the cabinet over the stove.

In contrast, to make sense of Mary's answer in (12), one must assume that Mary has just seen something or figured something out from which she concludes that the sugar is in the cabinet over the fridge. Perhaps Mary has seen a telltale trail of white particles, or perhaps she has realized that only one cabinet remains in the kitchen which Ann has not ruled out. This assumption has repercussions throughout the scenario. For if Mary is finding evidence and drawing inferences in this way, then most likely, she is searching with Ann for the sugar, somewhere where neither frequents. And if Ann asks her question under such circumstances, then Mary must just have made her discovery clear, perhaps by saying "Aha!"

The implicatures of Bill's response are also quite different from those of Mary's. Bill's utterance carries an implicature that Bill does not know whether the world is normal: perhaps he cannot remember whether this is one of the days he accidentally put the sugar over the stove. Indeed, suppose John subsequently looks over the fridge and finds no sugar. If this happened, could John then reproach Bill for the bad directions? My intuitions say no. Bill could respond that the sugar does indeed normally go in the cabinet over the fridge, that this instance is not normal, and that he warned John that this might be the case.

To make her reply in (12), meanwhile, Mary must believe that the sugar is in the cabinet over the fridge. There is no suggestion that the sugar will be there only if things turn out normally. For example, should Ann open the cabinet and find nothing, Mary would be proved wrong. Mary could not then say, "I was correct to infer that the sugar was in there, but this is just an abnormal circumstance."

Now, the pattern of accommodation and implicature for (11) is precisely what is predicted from the assumption that Bill's use of *should* semantically marks the information he presents as tentative or defeasible. On this assumption, the Maxim of Quality requires that Bill have evidence which justifies his tentative assessment: so Bill must know what is normally the case (and no more). Moreover, Grice's Maxim of Quantity allows the conventional implicature that Bill lacked information to make a more definite assertion. Otherwise, Bill ought to have used a stronger and briefer indicative statement

Take Kratzer's theory of modal bases and ordering sources in particular. Bill's answer is fleshed out as in (13).

(13) To any world compatible with my knowledge, there corresponds a more normal world in which the sugar is over the fridge.

The modal base is epistemic; the ordering source is what is normal. Thus to have grounds for (13), Bill must be able to judge the normality of worlds in the relevant respect. But by relativizing his claim that the sugar is in the fridge to only the most normal epistemically possible worlds, Bill implicates that he was not warranted in making a parallel assertion without the restriction.

On a semantic account like Kratzer's, the meaning *must* receives, and hence its conversational pragmatics, is essentially as described for *should*. This is not correct: *must* triggers different accommodations and implicatures from *should*. These effects depend on the fact that *must* presents a conclusion by referring to the evidence from which it is made.

Thus, consider the alternative paraphrase for Mary's response in (12):

(14) Some particular collection of facts A, salient in the common ground, provide a decisive reson to adopt the belief that the sugar is in the cupboard over the fridge.

The argument A gives Mary's utterance a special coherence relation with context, and forces the additional accommodation discussed above. For, felicitous reference to this A requires that there be some body of facts that Mary thinks Ann can recover that provide a reason for her conclusion, such as the trail of sugar which Ann could see, or the exhaustion of the search which she too could deduce. Meanwhile, the fact that (14) presents A as a justified argument that the sugar is over the fridge means that Mary is committing herself to the truth of its conclusion, as we found. The vague implicature of tentativeness for Mary's utterance can be explained by this paraphrase as well. Mary's position in asserting **must** p is weaker than that in asserting p, because it tips her hand: the argument A explicitly outlines a network of facts such that Mary will give up p if she is led to give up any of these facts. If Mary asserts p, in contrast, the assertion can only be challenged by attacking p itself.

The interpretation of (12) provides complementary evidence to the evidence presented earlier. In this discourse, too, we find that epistemic *must* differs from the indicative not in terms of the logical properties of the conclusion it presents, but in terms of the way it presents that conclusion.

5. Argument Systems

The examples in sections 3 and 4 show that formalizing this account depends on a model of reasoning in which defeasible conclusions may be adopted on the basis of packages of evidence natural enough that speakers can easily refer to them in conversation. Simari and Loui's argument systems provide such a framework. For Simari and Loui [16], arguments are "*prima facie* proofs that may make use of assertions that one sentence is (defeasible) reason for another (126)." Simari and Loui show how to use arguments to model the process of reasoning to a conclusion through a formal definition of JUSTIFICATION of arguments. In general, arguments only provide tentative support for propositions. However, those arguments with no effective counterarguments do justify the facts they support.

The Artificial Intelligence literature offers many other formalisms for defeasible reasoning, but argument systems are particularly appropriate for characterizing conclusions in a computational semantics for two reasons. (Similar reasons prompt the use of arguments for plan recognition in [4].)

The first is efficiency. Simari and Loui show how standard logic programming techniques and mathematical results about the possibility of finding conflicting arguments can be combined to enable relatively efficient searches that determine when an argument is justified in a particular context.

The second is separation of the local considerations that provide evidence from the global considerations that justify evidence. This differs from many default reasoning systems, such as McCarthy's circumscription [13] or Reiter's default logic [15]. In circumscription, a deduction by a defeasible rule proceeds only when abnormality predicates, minimized across the entire knowledge base, take on appropriate values. In default logic, the use of a defeasible rule depends on a demonstration that some fact is consistent with the knowledge base as a whole. Thus, these frameworks have only GLOBAL justifications for beliefs. Such global justifications do not seem perspicuous: linguistically, the data provides no reason to suppose that people REFER to the absence of conflicting evidence when presenting a conclusion. In argument systems, the argument localizes the reason for belief, in an argument, and recasts global considerations as properties of the argument. Thus, in argument systems, *must* can refer to the argument in an intuitive way and can assert that the argument is justified.

(Note, by the way, that while an assumption-based truth maintenance system also keeps track of the arguments that support conclusions, it reasons defeasibly by making and later retracting possibly incorrect assumptions rather than by representing justifications themselves as defeasible [2].)

The following definitions outline the development of argument systems in [16] (128–143). The formalization starts from a context κ consisting of a set K of established propositions consisting of ground formulas K_C and (logical) rules K_N and an additional set Δ of (defeasible) rules. The definition of arguments captures the notion of *prima facie* proof: an argument consists of some minimal collection of applications of defeasible rules which can be consistently added to K to obtain some conclusion. An argument's subarguments record ALL the conclusions that can be derived from a subset of the argument's premises and the logical rules in the context—not just those subarguments used in deriving the argument's conclusion.

Definition 1 *A set T of instantiations of elements of Δ is an* ARGUMENT *for a ground formula h ($\langle T, h\rangle_\kappa$) if and only if: (1) $K \cup T \vdash h$; (2) $K \cup T \nvdash \bot$; and (3) for no $T' \subset T, K \cup T' \vdash h$.*

Definition 2 *$\langle S, j\rangle_\kappa$ is a* SUBARGUMENT *of $\langle T, h\rangle_\kappa$ if and only if $S \subseteq T$.*

At any point, different pieces of evidence and different chains of reasoning may lead to different, possibly incompatible conclusions. One argument COUNTERARGUES another when the conclusion of the one argument contradicts an intermediate result of the other.

Definition 3 *$\langle T_1, h_1\rangle_\kappa$* COUNTERARGUES *$\langle T_2, h_2\rangle_\kappa$* AT *$\langle T, h\rangle_\kappa$ if and only if $\langle T, h\rangle_\kappa$ is a subargument of $\langle T_2, h_2\rangle_\kappa$ and $K \cup \{h, h_1\} \vdash \bot$.*

In such cases of conflict, the MORE SPECIFIC ARGUMENT, which takes more of the particulars of the context into consideration, is preferable and is said to DEFEAT the less specific argument. The definition of specificity measures dependence on particulars using the intuition that a more general argument $\langle T_2, h_2\rangle_\kappa$ yields its conclusion from any piece of evidence e for which the more specific argument $\langle T_1, h_1\rangle_\kappa$ yields its conclusion, but not vice versa.

Definition 4 *$\langle T_1, h_1\rangle_\kappa$ is* MORE SPECIFIC *than $\langle T_2, h_2\rangle_\kappa$ if and only if: (1) for all ground formulas e, if $K_N \cup \{e\} \cup T_1 \vdash h_1$ but $K_N \cup \{e\} \nvdash h_1$, then $K_N \cup \{e\} \cup T_2 \vdash h_2$; and (2) there is some ground e such that $K_N \cup \{e\} \cup T_2 \vdash h_2$, $K_N \cup \{e\} \cup T_1 \nvdash h_1$, and $K_N \cup \{e\} \nvdash h_2$.*

Definition 5 *$\langle T_1, h_1\rangle_\kappa$* DEFEATS *$\langle T_2, h_2\rangle_\kappa$ if $\langle T_1, h_1\rangle_\kappa$ counterargues $\langle T_2, h_2\rangle_\kappa$ at $\langle T, h\rangle_\kappa$ and $\langle T_1, h_1\rangle_\kappa$ is more specific than $\langle T, h\rangle_\kappa$.*

An argument is justified whenever it has no counterarguments which are not themselves defeated. This notion is captured in two stages, using an inductive definition of supporting and interfering arguments. This induction describes the provisional status of arguments as warranting their conclusions (supporting arguments) or as casting doubt on other arguments (interfering arguments) as the effects of defeat among arguments propagate. An argument is justified if, after a certain point in this induction, no further evidence against it comes to light.

Definition 6 *All arguments are level 0 supporting and interfering arguments.*

- *An argument $\langle T_1, h_1\rangle_\kappa$ is a level $(n+1)$ supporting argument if and only if no level n interfering argument counterargues it at any of its subarguments.*
- *An argument $\langle T_1, h_1\rangle_\kappa$ is a level $(n+1)$ interfering argument if there is no level n interfering argument which defeats it.*

Definition 7 *An argument $\langle T, h\rangle_\kappa$* JUSTIFIES *$h$ in κ if and only if there is some m such that for all $n \geq m$, $\langle T, h\rangle_\kappa$ is a level n supporting argument. h is justified in κ if some $\langle T, h\rangle_\kappa$ justifies it in κ.*

I write $\kappa \models \langle T, h\rangle_\kappa$ if $\langle T, h\rangle_\kappa$ justifies h in κ and $\kappa \models h$ if h is justified in κ.

I have constructed a program to construct arguments and establish their justification, similar to the implementation described by Simari and Loui. First, the deductive closure of a database in minimal horn logic is computed by forward-chaining; using the derived proofs of $\bot$, the minimal arguments are isolated with which each counterargument is in conflict. This smaller collection of arguments is then evaluated according to an operationalization of the definition of support and interference appropriate because this database contains only finitely many arguments, as follows. The arguments are traversed depth first: an argument is marked directly as defeated, interfering, or justified according to the recursively computed marking just of those arguments which potentially defeat or interfere with it. One visit is enough to compute the marking of each argument. With this labelling made, remaining arguments in the database inherit the weakest labelling of any argument used in their construction, because the status of all counterarguments has been determined.

5.1. Linguistic Application

If we adopt some context or database κ as the representation of the information state of an individual or conversation, and assume a translation of an indicative sentence S into a logical formula $F(S)$, we can use argument systems to define the truth of a sentence as follows:

(15) $\underline{S}$ is true in κ if and only if $\kappa \models F(S)$.

We can, moreover, model indicative epistemic conditionals involving extension of information using the Ramsey Test [6].

(16) $\underline{\text{If } S\text{, then } T}$ is true in $\kappa = \langle K_C, K_N, \Delta\rangle$ if and only if $\langle K_C \cup \{F(S)\}, K_N, \Delta\rangle \models F(T)$.

Most interestingly, we can give this definition for an utterance of **must** S that refers to an argument A:

(17) $\underline{\text{Must } S(A)}$ is true in κ if and only if $\kappa \models \langle A, F(S)\rangle_\kappa$.

A concrete example will help to illustrate how these definitions fit together. Consider the pyrotechnic knowledge base Δ_0 of defeasible rules given in Figure 1. The first two clauses are intended to represent (schematically) that a struck match tends to have been lit; but that a wet, struck match tends not to have been lit. The other two clauses describe two ways things are normally hot: if they have been plunged into boiling water, or if they have been set on fire.

(18) $\texttt{match(x)} \wedge \texttt{strike(x)} > \texttt{lit(x)}$

(19) $\texttt{match(x)} \wedge \texttt{strike(x)} \wedge \texttt{wet(x)} > \neg\, \texttt{lit(x)}$

(20) $\texttt{boiled(x)} > \texttt{wet(x)} \wedge \texttt{hot(x)}$

(21) $\texttt{lit(x)} > \texttt{hot(x)}$

Figure 1: Δ_0.

Consider the resolution of the arguments available in

(22) $\kappa_0 = \langle\{\texttt{match(m1)}, \texttt{struck(m1)}\}, \{\}, \Delta_0\rangle$

in which the rules of Figure 1 apply to a single struck match. In this context, forward-chaining constructs two arguments:

(23) $A_1 = \langle\{(18)\}, \texttt{lit(m1)}\rangle_{\kappa_0}$
$A_2 = \langle\{(18), (21)\}, \texttt{hot(m1)}\rangle_{\kappa_0}$.

Since no counterarguments result, both are marked justified. Accordingly, in the context

(24) $\kappa_1 = \langle\{\texttt{match(m1)}\}, \{\}, \Delta_0\rangle$

sentence (25) is true.

(25) If the match was struck, it lit.

To interpret this, we simply add $\texttt{struck(m1)}$ to κ_1, obtaining κ_0, forward-chain to construct arguments, and test whether $\kappa_0 \models \texttt{lit(m1)}$. But it does, since $\kappa_0 \models A_1$.

Of course, (26) is also true in κ_1.

(26) If the match was struck, it must [by A_1] have lit.

The reason is exactly the same.

Now turn to a context in which it is known that the match was struck but that it was wet.

(27) $\kappa_2 = \langle\{\ \texttt{match(m1)}, \texttt{struck(m1)}, \texttt{wet(m1)}\}, \{\}, \Delta_0\rangle$

Again in κ_2, we derive A_1 and A_2. However, another argument will appear: $A_3 = \langle\{(19)\}, \neg\texttt{lit(m1)}\rangle$. A_3 and A_1 together derive the only contradiction in this context; since A_3 is more specific than A_1, A_3 defeats A_1 and not vice versa. In fact, A_3 has no defeaters. Thus, A_3 is justified, while A_2 and its subverted subargument A_1 are not. (This characterization accords with the inductive definition: only A_3 is a level 1 interfering argument, since it defeats both A_1 and A_2. Thereafter, A_3 becomes the sole supporting argument, as its counterarguments have been defeated.)

Thus, back in κ_1 both of these sentences are true:

(28) If the match was struck when wet, it { did not light, must [by A_3] not have lit }.

The contrast between (25) and (28), reminiscent of the Samaritan paradox discussed by Kratzer [11], illustrates the defeasibility of inference in argument systems. Defeasibility on this view accounts for the patterns of inference from which Kratzer motivates the relativization of modal meanings to an ordering source.

6. **must** S versus S

As the discussion of (25) and (28) shows, on this account, **must** S CAN BE truth-conditionally equivalent to S, in the following sense. When and only when S can be truly utterred, there is some argument A that justifies it. Provided a speaker felicitously intends to refer to this argument in uttering **must** S, that utterance too is truthful.

However, **must** S may be false when S is true, if the modal utterance refers to an unjustified argument. Here is an example of this: take the context

(29) $\kappa_3 = \langle\{\ \texttt{match(m1)}, \texttt{struck(m1)}, \texttt{boiled(m1)}\}, \{\}, \Delta_0\rangle$

(30) $B_1 = \langle\{(18)\}, \texttt{lit(m1)}\rangle_{\kappa_3}$

(31) $B_2 = \langle\{(18), (21)\}, \texttt{hot(m1)}\rangle_{\kappa_3}$

(32) $B_3 = \langle\{(20), (19)\}, \neg\texttt{lit(m1)}\rangle_{\kappa_3}$

(33) $B_4 = \langle\{(20)\}, \texttt{hot(m1)}\rangle_{\kappa_3}$

Figure 2: The arguments of κ_3.

In κ_3 we have the arguments shown in Figure 2. Here B_3 counterargues and defeats B_1 and B_2. (B_3 depending on the premises that the match be boiled, as well as struck, is more specific than the two that fail to take the boiling into account.) But B_4 is not counterargued at all. Thus, the justified arguments are B_3 and B_4.

Because of B_4, the following sentences are both true in κ_3:

(34) The match { is, must [by B_4] be } hot.

But the following sentence is not:

(35) The match must [by B_2] be hot.

B_2 is not justified in κ_3, even though its conclusion is. A sense of what this means can be provided by the dialogue in (36), which assumes the same context, and which shows how a speaker might signal the dependence of a modal on this particular argument, and how another speaker might respond.

(36) One: The match was struck, so it must be hot.
Two: Well, no. It is hot because it was boiled. It didn't light.

In sum, **must** S, unlike S, inherits the strength of the argument by which the speaker justifies it. This fact is crucial in explaining the kinds of inferences we found presented with *must* in section 3. Because of its reference to an argument, epistemic *must* is appropriate in two kinds of circumstances where it is important for a speaker to signal the dependence of a claim on particular evidence.

The first case arises when the evidence is particularly forceful or surprising. Then, referring to the argument ensures that the addressee recovers the argument and considers the argument for himself. This may make the modal assertion stronger or more convincing than the corresponding indicative assertion, where the addressee might not consider this evidence, or might not realize that the speaker has considered it. Hence the prevalent use of this kind of *must* in establishing the coherence of mathematical and other technical, intricate demonstrations. (Recall example (6).)

The second case arises when the speaker intends to make manifest the weakness of his evidence. The speaker uses *must* to signal what evidence he has, and thus to prevent the hearer from incorrectly inferring that the speaker might have other, stronger evidence for his proposition. This use of modality seems weaker than the corresponding indicative assertion. For, as alluded to earlier and as in seen in (36), referring to such an argument leaves the speaker's claim more vulnerable to attack than it would otherwise be.

It is important to note that in either case, the speaker is using *must* to present the best evidence available to him for a proposition: evidence on the basis of which he would adopt the proposition into his beliefs. The motivations given for using *must* above explain variation in force not because *must* marks reasoning as in any way defective, but because the underlying motivations for using *must* reflect how the speaker takes into account his own capacity for ignorance, error and oversight and the limited inferential resources of his addressee. If conclusions were presented as defective, we would as shown earlier obtain the pattern associated with the word *should*.

7. On Referring to Arguments

To complete this account, I must explain precisely how reference to an argument is achieved. In a nutshell, speakers refer to arguments the same way they refer to anything else: they present enough constraints on the object in question so that all objects except the intended one can be ruled out either as incompatible with the constraints or as too remote to have been intended. This is how, for example, a definite noun phrase refers to the most salient discourse entity that matches the description's content. (See Sperber and Wilson [17] or Grosz and Sidner [8] for more details.)

In the case of epistemic *must*, constraints typically specify a fact that CONTRIBUTES to the reference argument.

Definition 8 *A fact f* CONTRIBUTES *to* $\langle T, h\rangle_\kappa$ *in* $\kappa = \langle K_C, K_N, \Delta\rangle$ *if and only if f is an instantiation of an element of* K_C*,* K_N *or* Δ *which in turn can be instantiated to an element of some least set* $\Gamma \subset K_C \cup K_N \cup T$ *such that* $\Gamma \vdash h$*.*

In other words, a fact contributes to an argument if it is necessary to make that argument go through.

Such contributing facts can be explicitly supplied using several interesting constructions. The locution *must p the way q*, illustrated in (37), is common in colloquial speech.

(37) It must have been hurt the way it was limping.

In (37), we might assume that *the way q* is a syntactic argument of *must* that supplies a characterization of the (reasoning) argument *must* refers to. On this analysis, the locution receives the paraphrase that *p* is the conclusion of a (salient) justified argument containing *q that way* as a premise. For (37), this gives:

(38) The fact that it was hurt is the conclusion of a justified argument containing as a premise the fact that it was limping how it was.

Note in comparison the difficulty of explaining the role of clauses like *the way it was limping* without arguments. For example, on the possible-worlds account, the clause should already be part of the modal base, since it is known; but it is unclear how to incorporate this single proposition meaningfully into the ordering source.

In the Brown corpus, clauses linked with modal assertions, using *because* or using colons or semicolons, often seem to play a similar role to the *way*-clause in speech (cf. (6)):

(39) A number of [places on the list] must have fallen into disfavor; they were struck out with remarks in red ink, denouncing both the cooks and the management.

(40) I must have dozed off, because I came to with a start at the sound of voices.

Most commonly in the Brown Corpus, epistemic *must* appears in a context where no such constraint is presented, but the reference argument still can be determined, because its premises are made salient by close mention nearby (cf. (8)):

(41) William Beebe reports 26 inches and 2.4 ounces (this snake must have been emaciated) for the length and the weight of a young anaconda from British Guiana .

In fact, the descriptive content of sentences seems frequently to have been enriched to make the right information salient, as below where the length of the ride probably explains why he got hungry on the ride more than it distinguishes this ride from others (cf. (7)):

(42) That long ride the four of you took must've given him a good appetite .

As we saw earlier in Mary's reply to Ann, it may be necessary to accommodate a salient argument with *must*. This accommodation may in fact be intended, as in (43), where the writer seems to refer to, but not to specify, some aspect of Garth's appearance from which his age can be judged.

(43) He disliked Garth on sight, conservative clothes and haircut, smile a shade too earnestly boyish for a man who must be well into his thirties, handclasp too consciously quick and firm.

8. Conclusion

I have argued that epistemic *must* allows speakers to present conclusions by referring to the arguments that justify them. Support for this claim derives from characteristic patterns of use that sentences containing *must* display. Particularly important are these sentences' implicatures and their need for coherence with surrounding discourse, as seen particularly through the accommodation they trigger. You might say that on this view, the meat of the job of *must* in presenting a conclusion lies in what it means to PRESENT, not what it is to CONCLUDE.

The importance of presentation to the meaning of *must* may at first seem surprising. It is by no means unparalleled, however. The perfect, for example, is quite similar: speakers use the English perfect when it serves their communicative intentions to make the connection between the result state of an action and a later situation particularly clear, not when they describe an action with specific causal properties. For limited agents like ourselves, making clear how and why an assertion is made can be important as making the assertion itself [19]. Epistemic *must* is one of many words we have to help do this.

References

[1] Jennifer Coates. *The Semantics of the Modal Auxiliaries.* Croom Helm, Beckenham, Kent, 1983.

[2] Johan de Kleer. Extending the ATMS. *Artificial Intelligence*, 28:163–196, 1986.

[3] M. Ehrman. *The Meanings of the Modals in Present-day American English.* Mouton, The Hague, 1966.

[4] George Ferguson and James F. Allen. Arguing about plans: Plan representation and reasoning for mixed-initiative planning. In Kristian Hammond, editor, *Proceedings of the Second International Conference on A.I. Planning Systems*, pages 43–48, 1994.

[5] W. Nelson Francis. *A Standard Sample of Present-day English for use with Digital Computers. Report to the U.S. Office of Education on Cooperative Research Project No. E–007.* Brown University, Providence, 1964.

[6] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* MIT Press, Cambridge, MA, 1988.

[7] H. P. Grice. Logic and conversation. In P. Cole and J. Morgan, editors, *Syntax and Semantics III: Speech Acts*, pages 41–58. Academic Press, New York, 1975.

[8] Barbara Grosz and Candace Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12:175–204, 1986.

[9] Gilbert Harman. *Change in View: Principles of Reasoning.* MIT Press, Cambridge, MA, 1986.

[10] Angelika Kratzer. The notional category of modality. In H. J. Eikmeyer and H. Rieser, editors, *Words, Worlds and Contexts*, pages 38–74. de Gruyter, Berlin, 1981.

[11] Angelika Kratzer. Modality. In A. von Stechow and D. Wunderlich, editors, *Semantics: An international handbook of contemporary research*, pages 639–650. de Gruyter, Berlin, 1991.

[12] Fred Landman. Data Semantics: An epistemic theory of partial objects, in *Towards a Theory of Information: The Status of Partial Objects in Semantics*, pages 1–96. Foris, Dordrecht, 1986.

[13] John McCarthy. Applications of circumscription to formalizing commonsense knowledge. *Artificial Intelligence*, 28:89–116, 1986.

[14] F. R. Palmer. *Mood and Modality.* Cambridge Univeristy Press, Cambridge, 1986.

[15] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 12:81–132, 1980.

[16] Guillermo R. Simari and Ronald P. Loui. A mathematical treatment of defeasible reasoning and its implementation. *Artificial Intelligence*, 53:125–157, 1992.

[17] Dan Sperber and Deirdre Wilson. *Relevance: Communication and Congition.* Harvard University Press, Cambridge MA, 1986.

[18] Frank Veltman. Data semantics. In Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, editors, *Truth, Interpretation and Information.* Foris, Dordrecht, 1981.

[19] Lyn Walker. *Informational Redundancy and Resource Bounds in Dialogue.* PhD thesis, University of Pennsylvania, 1993.

[20] Erich F. Woisetschlaeger. *A Semantic Theory of the English Auxiliary System.* Garland, New York, 1985.