

DE-MAXIM-IZING QUALITY

CHRISTINE GUNLOGSON
University of Rochester

1 Introduction

Quality has long been recognized as having a unique status among Grice's conversational maxims, shown in 1–4 (Grice, 1989b):

- (1) Quality: Try to make your contribution one that is true.
 - a. Do not say what you believe to be false.
 - b. Do not say that for which you lack adequate evidence.
- (2) Quantity
 - a. Make your contribution as informative as is required (for the current purposes of the exchange).
 - b. Do not make your contribution more informative than is required.
- (3) Relation: Be relevant.
- (4) Manner: Be perspicuous.

Grice himself was the first to remark on the exceptional nature of Quality (Grice, 1989b:27):

In modern Gricean approaches, Quality is invoked in several ways. It is standardly assumed to underlie inferences about the speaker's propositional attitude toward the content of her contribution. If the speaker is meeting the demands of Quality—as interlocutors normally presume—then addressees can infer that the speaker believes what she is saying and (believes she) has sufficient evidence for it.¹ It does not follow that the addressee necessarily believes what the speaker says is true, but acceptance is often assumed as a simplification. The presumption that the speaker intends to be conforming to Quality is also crucial to the logic of scalar implicatures, where the necessity of satisfying Quality serves an upper bound on the amount of information a speaker can cooperatively provide.

Quality has received relatively little attention in comparison to Relation and Quantity. Formal dialogue models often stipulate truth-telling as a mandatory rule of the discourse, or limit attention to dialogues where it is observed — a useful idealization for many purposes, but one that renders the operation of Quality nearly invisible.

In this paper I review and add to the ways in which Quality is known to differ from the other maxims. These differences are not superficial but represent important properties which must eventually be accounted for under anyone's theory. The claim is that to regard Quality as a routinely-observed maxim is to obscure some of the real puzzles that surround its operation and the inferences it gives rise to.

I lead off in Section 2 with several non-standard features of Quality-based inferences: their noncancelability, their origin (in the usual case) in satisfaction of the maxim rather than violation, and their capacity for being suspended. Section 3 goes on to consider how the domain of Quality compares to that of the other maxims and raises prospects of generalization to implicit (unuttered) content and to non-assertive and/or non-declarative utterances. The cumulation of the differences discussed makes what I believe to be a strong case for 'de-maxim-izing' Quality while retaining it as a fundamental principle of a different sort. Precisely what this 'different sort' is, however, remains unclear. In Section 4 I outline some of the issues that arise in attempting to formulate the principle of Quality as a condition on context. In service of the overall goal of this paper, which is to lay out the territory that any full-fledged implementation of Quality should be responsible for, Section 5 summarizes the properties of Quality identified throughout the preceding discussion.

2 Qualities of Quality

The maxim of Quality, like the other maxims, and the overarching Cooperative Principle that they explicate, is stated as a set of guidelines for *speakers*. Given that language users are, as a rule, both speakers and interpreters, guidelines that speakers reliably follow will constitute reliable guidelines for interpreters as well, triggering inferences based on conversational and linguistic choices made by speakers presumed to be cooperative.

1

2.1 Non-cancelability

The primary inference that Quality is responsible for is the attribution of belief to the speaker, in the absence of counterindications. Treating the inference as a pragmatic, utterance-based implicature is a solution that many have found appealing, especially if the alternative is building an implicit propositional attitude into sentence meaning. This Quality-based implicature is sometimes said to explain the Moore's paradox effect exhibited in 5 (see, for instance, Martinich (1980); Gamut (1991); Huang (2007)).

- (5) Moore's paradox
 - a. The cat is outside, though Jack doesn't believe that it is.
 - b. The cat is outside, /though I don't believe that it is.

Under the implicature view, the 1st person expression of disbelief in 5b is anomalous because it contradicts the implicature arising from the first clause, namely that the speaker believes the cat is outside. No such contradiction occurs when someone other than the speaker is skeptical, as in 5a.

But 5b also illustrates the fact that inferences about belief arising from Quality lack one of the defining properties of conversational implicature, cancelability. This exceptional behavior of Quality 'implicatures' has not gone unnoticed; see remarks in Levinson (1983) and Horn (1984), for instance.

For comparison, cancelations of Quantity-based implicatures are shown below. The second clause in 6 cancels the *not all* implicature associated with *most*; 7a does the same for the exclusive sense reading of *or*:

- (6) Most of the students were there, in fact they all were.
- (7) May or Joanna will be there,
 - a. and perhaps both. [cancelation]
 - b. /and perhaps neither. [contradiction]

The contradictory flavor of 7b indicates that the first and second clauses are at odds with each other, diagnosing an entailment relation. The conclusion is that the first clause entails the negation of 7b, i.e., entails that at least one of the set *fMay*;

according to Grice, and the single example cited involves flouting and no other maxims.³ Violation of the maxim of Manner is likewise discussed only in connection with flouting.

If the limited scope for “maxim clash” suggested above is accurate, it is better cast in a more particular light as the kind of predictable interaction between Quality and the 1st submaxim

of Relation-based implicatures. As mentioned earlier, the presumption that a contribution will conform to the maxim of Relation is so strong that when something apparently irrelevant is said, the machinery of inference sets to work to discover the way(s) in which it is relevant after all, hence not a true violation. By contrast, effects Grice attributes to flouting of (the first maxim of) Quality, such as sarcasm, do not involve reconciling the mismatch between what the speaker says and what the speaker can be presumed to believe. Instead they require recognizing the disparity and drawing inferences from the speaker's choice to present it, overriding the inferences of speaker belief that normally obtain.

2.3 Suspension

As we saw in 2.1, Quality-based inferences about speaker belief are not individually cancelable.

Furthermore, within non-epistemic contexts, ordinary sorts of implicatures still go through, despite the suspension of Quality-based inferences of speaker belief. If, for example, the narrative features a genie who tells the hero he gets three wishes, we take that to mean exactly three; the use of *some* by a narrator or character will still implicate *not all*; and in 11, Green's implicated *no* to Hornsby's proposal is clear even though they are invented characters (*The Night Gardener*, (Pelecanos, 2007):60):

- (11) "Anybody up for a beer?" said Hornsby. "I'll let y'all buy me one."
 "I got practice," said Green...

It seems that recognizing an implicature, even a particularized implicature as in 11, not only doesn't require being a participant in the discourse that produced it, it doesn't require that the implicating speaker even exist, let alone have communicative intentions of the appropriate sort.

Much remains to be said on this topic, but leaving fictional implicature for further research, I turn next to consideration of Quality inferences from unuttered contributions.

3 The domain of Quality

The maxim of Quality, repeated below for reference, makes use of the term 'contribution' and prohibits 'saying' what is false or insufficiently supported.

- (1) Quality: Try to make your contribution one that is true.
 a. Do not say what you believe to be false.
 b. Do not say that for which you lack sufficient evidence.

What does the contribution include, and what counts as saying? Interpreting 'say' in Grice's 'favored sense' (Grice, 1989d:25), the answer is quite clear: Quality is stated in a way that narrows its applicability to conventional linguistic content, something like the literal meaning of the sentence. Attempts to formalize the maxims, such as Gazdar (1979), Groenendijk (1999) have tended to implement Quality as a condition on assertive utterances of declarative sentences, compatible with this narrow interpretation. Gazdar's formulation provides a usefully explicit instance. Noting the exceptionality of Quality as a maxim, he implements it instead as a felicity condition on assertoric utterances. The informal version of the condition is given below (corresponding to Gazdar's (32)); $K\phi$ is short for *the speaker knows that ϕ* :

- (12) For any declarative sentence ϕ , **assertion** of ϕ commits the speaker to $K\phi$.

Gazdar's condition invokes both sentence type (declarative) and speech act (assertion), restricting its application to a subset of utterances of declaratives. While few analysts are as explicit about assumptions, the implicit restriction of assertions/statements to declarative utterances, and conversely, of declarative utterances to assertion, is widely adopted in practice.

In this section I will raise some questions about this narrow conception of what Quality 'applies to'. The first concern, addressed in Section 3.1, relates to the status of implicit (unuttered) contributions to the discourse with respect to Quality considerations. The second is the issue of whether, and how, to extend Quality-like principles to other kinds of sentence types and functional categories, assuming that they also lead to inferences about the speaker's state. I will turn to that problem in Section 3.2.

3.1 Implicit content

In this section I suggest that Quality, unlike the other maxims, can usefully be broadened to include implicit — unuttered — content. This is a departure from Grice's treatment. Implicatures and (accommodated) presuppositions are two prominent examples of such implicit contributions. At issue is whether the sort of Quality-based attributions of speaker attitude that accompany asserted content also attach to implicated and presupposed content. That is, do we normally infer that speakers believe and have evidence for what they are implicating or presupposing, just as we do for the content they are explicitly uttering? I believe the answer is yes. If that is correct, the standard formulation of Quality as applying only to the content of utterances, narrowly construed, should be reconsidered. It is possible, of course, that inferences of speaker belief associated with implicit content are not due to the Quality-based presumption of truth-telling, but have some different source. But Quality is the first place to look on grounds of parsimony alone. Let us start with presuppositions.

3.1.1 Presuppositions

I take it to be uncontroversial that we do generally take speakers believe and have evidence for what they presuppose, or at least that if they believe what they are asserting, they also believe what they presuppose in asserting. To make this point more concrete, consider the examples in 13. If a speaker asserts 13a, he cannot deny that Jack has a cat without contradiction. A similar point holds for belief attribution. An addressee who infers that the speaker of 13a believes Jack's cat is

would not be required. Depending on one's view of presupposition, there are other possibilities. If presuppositions are taken to be part of the literal content of the sentence, associated with their lexical trigger, belief attribution could presumably be covered under a suitable implementation of Quality as it stands. A similar result (for present purposes) would follow from taking a purely pragmatic view of presupposition like that of Stalnaker (1978), where presuppositions have the status of background beliefs the speaker must be assumed to hold if she believes what she asserts. Attribution of speaker belief for presuppositions follows from the attribution for asserted content.

All three non-Quality maxims are sensitive to the content of what is said, often at the level of lexical choice, for Quantity at least. The implicature-generating mechanism relies on evaluation of what was said against some notion of what a cooperative speaker, acting in accordance with the maxims, *ought* to offer. The actual content that the speaker offers must be assessed in this comparison.

Quality-based inferences about speaker belief spring from a different source. The inference that the speaker believes the content she is asserting does not depend on evaluating that content against the standard of what a cooperative speaker would assert. In fact, the content of the utterance matters for these Quality inferences only to a limited extent: flouting must be recognized as such, and there must be some sensitivity to sentence type/speech act⁴ to restrict applicability to content that can be evaluated in terms of truth and evidence. Given a minimally appropriate utterance, i.e. one whose content can plausibly be taken as a belief of the speaker's, the attributive inference is the default, regardless of content.

If this characterization of Quality's operation is accurate, there is nothing to prevent broadening its scope to apply to implicatures as well as assertions. The result of extending the domain of Quality would be to activate the default attribution of speaker belief for the implicated as well as the asserted content of a speaker's 'contribution'. I believe this is a desirable result — that we do, by default, take others to believe what they communicate via implicature, and understand a speaker to be misleading when she implicates what she does not believe and have evidence for. Possible objections to this claim will be considered shortly.

One way to accomplish the extension of Quality's domain would be to make it a default contextual condition, applying to all (propositional) additions to the context, explicit or implicit. I will take up that idea in Section 4. For now, I will make a simple adjustment to the wording of Quality, at the Gricean level of informality, substituting *communicate* for *say*.

- (15) Quality [revised]: Try to make your **communicative** contribution one that is true.
- a. Do not **communicate** what you believe to be false.
 - b. Do not **communicate** that for which you lack sufficient evidence.

Setting implicatures aside for a moment, there seems to be independent justification for extending Quality to communicated content. Consider communicative gestures such as nodding and pointing, which do not ordinarily count as 'saying' or 'uttering' something. Surely it is just as deceitful to give a false answer by nodding as it is to utter *yes* — it is the conflict between the facts and the affirmative response that is the problem, regardless of how the affirmation is conveyed. (For that matter, it is not entirely clear how *yes*, *no*, and other response particles come by the appropriate 'literal content' to generate Quality inferences.) Similarly, an unsupported accusation delivered via pointing (to someone in a line-up, say) is as deserving of censure as uttering, "That's

utterances. The fact that they are not uttered is not incidental but definitional. A nod can count as

3.2 Beyond statements

Quality is the only maxim to mention truth and falsity (emphasis added below):

- (1) Quality: Try to make your contribution one that is **true**.
 - a. Do not say what you believe to be **false**.
 - b. Do not say that for which you lack adequate evidence.

The other maxims make reference more generally to the speaker's 'contribution' without directly invoking truth values and attitudes. The stated goal of making a *true* contribution means that Quality in its original form does not extend straightforwardly beyond the category of statements/declaratives — that subset of utterances that can be evaluated in truth-conditional terms. I suspect that the problems discussed here arise in some form for the other principles as well, particularly for Quantity, where the notion of relative informedness lends itself to modeling in terms of entailment relations. Nevertheless, I will continue to concentrate on Quality, which in Grice's rendering is exceptionally explicit about the nature of the contribution required.

The restriction of Quality to truth-conditional content raises a number of complex issues. I single out for discussion here the matter of inferences about speaker attitudes arising from utterances falling outside that domain. Do we need multiple versions of Quality, each applying to a distinct kind of utterance? And if so, what are the appropriate utterance categories? One point to be kept in mind is that once we accept that the domain of (the original) Quality is restricted, the necessity of classifying utterances by type is unavoidable. This is true even if we do not adopt multiple versions of Quality. At a minimum there needs to be a binary distinction between utterances to which the original version is applicable and the 'other' category of utterances to which it is not. Thus, we must arrive at some method of categorization even to implement the basic version of Quality.

To begin the investigation, compare the inferences that would naturally follow from utterances of 16 and 17 in ordinary circumstances:

- (16) The server's down.
- (17) Is the server down?

We can take 16 to be the standard case of an assertive declarative, with the usual attribution of speaker belief arising from Quality. Inferences drawn from an utterance of 17 would be quite different. Normally the addressee would be justified in concluding that the speaker does not hold the belief that the server is down, in fact does not know whether that is the case, and does believe that the addressee may be able to supply a true answer. What is responsible for the routine inference of speaker ignorance attendant on an utterance of 17?

Similarly, the speaker's preference for a closed door will ordinarily follow as an inference from an utterance of 18. Inferences about speaker desires, however, are not readily explained by appeal to the truth-oriented standard of the original Quality maxim.

- (18) I prefer the door to be closed.
- (19) Close the door, please.

Note with regard to 18 – 19 that the inference supported in each case about the speaker's state (belief, ignorance, preference) is not presuppositional — that is, it does not have to be established in the context *before* the utterance. Rather, the inference follows from the utterance.

Taking note of the inapplicability of the Quality maxim to non-assertive utterances, Levinson (1983:105) suggests in passing that the truthfulness applies only to assertion and can be considered a subcase of a more general sincerity requirement for speech acts. Sincerity for questions, e.g., does not require speaking the truth – instead, the sincere questioner must truly desire an answer. The questioner cannot truly desire knowledge she already possesses. Therefore, presuming sincerity on the questioner's part goes hand in hand with presuming ignorance of the answer as well.

Along similar lines, Martinich (1980) criticizes the narrowness of Grice's version of Quality and proposes replacing it with a more general Authenticity requirement, shown in 20 (Martinich's B^h):

(20) Authenticity (Martinich, 1980)

Be authentic. That is, do not knowingly participate in a speech act for which the conditions for its successful and non-defective performance are not satisfied.

Martinich's idea is that Quality in its original form is a specialized instance of the general principle in 20, and thus applies to only to the assertive family of speech acts.

For concreteness, I offer an informal generalization of Quality from a functional point of view in 21, with hypothetical speech act categories italicized:

(21) Be sincere. (Quality)

- a. *Assert* only what you believe to be true and have evidence for.
- b. *Ask* only for information that you lack, want, and believe the addressee can supply in accordance with clause (a).
- c. *Issue a command* only to present an outcome that you wish the addressee to bring about.

I have shown only three clauses above, involving speech acts at a high level of generality. The appropriate number and granularity would depend upon the taxonomy of speech acts in whatever theory is adopted. In principle there could be a very large number of clauses, with sincerity conditions for all speech acts residing under the general Quality principle.

Levinson's and Martinich's proposals both suggest, without going into detail, that conditions on speech acts can account for the different kinds of inferences that arise from the canonical uses represented by 16, 17, and 19. Presumably the default assumption that speakers are sincere would lead to specific inferences for each speech act, as outlined above for questions. If a sincere questioner is ignorant of the answer, then the presumption of sincerity supports the inference that a speaker posing a question is ignorant. This suggestion is problematic, however, as can be seen most readily by taking an addressee's perspective. It is not much of a stretch to assume that a speaker embarking on a speech act is aware of the appropriate sincerity conditions for the act he has in mind. The same cannot be said for his addressee(s), who is not privy to the speaker's state and who therefore must identify the nature of the act from what is publicly available — the utterance itself. Context can be of assistance, to be sure. But as was pointed out in connection with the examples above, the context need not provide any insight into the speaker's state *before* the utterance. It is certainly possible for the utterance itself to be the sole or primary basis for inferences about the speaker's attitude and intentions.

The above remarks highlight what I take to be the most serious problem of a speech-act approach to generalizing Quality: the risk of circularity. The potential for circularity lies in

the assumed existence of some independent mechanism for identifying speech act category — independent, that is, of the general pragmatic principles subsumed by the Cooperative Principle, and independent of Quality in particular. I know of no such mechanism, and it seems quite likely that any method of categorizing speech acts will have to employ general pragmatic principles if it is to meet with any success. Searle (1975), in fact, is quite explicit about the appeal to Gricean reasoning in identifying indirect speech acts. But if the maxims are needed to identify the speech act, obviously they cannot be applied selectively in a way that requires a prior determination of what the speech act is.

There is an alternative system that avoids circularity, namely, categorization by *sentence type*. Following Sadock and Zwicky (1985), I use the term *sentence type* for the handful of major sentence classes distinguished formally at the root clause level, including declaratives, interrogatives, and imperatives. Because sentence type distinctions are rooted in formal properties rather than functional ones, no circularity arises. Recognition of a sentence as declarative, interrogative, etc. follows from the ordinary linguistic knowledge that distinguishes, e.g., *It's raining.* from *Is it raining?*

The sentence-type approach is in the spirit of Grice's own views, though he spoke of the effects of 'indicative mood' rather than sentence type. It is also consistent in a broad way with Searle's views as expressed in Searle (1975), and implicit in a number of other approaches. These approaches have in common that they (implicitly or explicitly) assign significance to the *kind* of sentence uttered, where that kind can be identified in terms of linguistic properties rather than function. I include in this group proposals that refer to actions such as 'stating', 'questioning', etc. without giving any other indication of how these actions are to be recognized. For instance, 'conversational postulates' proposed by Gordon and Lakoff (1975) for deriving indirect speech acts refer to *asserting* speaker-based sincerity conditions and *querying* hearer-based ones. But since the postulates themselves are used to identify a speech act, the actions of asserting and querying require grounding elsewhere.

The sentence-type categories mentioned in 21 are necessarily quite broad and are not expected to correspond to specific speech acts. This generality seems appropriate for a high-level guiding principle of the sort we expect Quality to be. Still, the revised generalization in 22 is oddly form-specific:

(22) Be sincere. (Quality)

- a. Utter a *declarative* sentence only to present what you believe to be true and have evidence for.
- b. Utter an *interrogative* sentence only to present information that you lack, want, and believe the addressee can supply in accordance with clause (a).
- c. Utter an *imperative* sentence only to present an outcome that you wish the addressee to bring about.

Note that, in contrast to 21, the number of clauses in 22 is a fixed small number, the number of major sentence types in the language.

Martinich cites declaratives functioning as questions, as in *That's a cat?*, as decisive evidence that linking the form of the sentence to speaker belief, as Grice suggests, is ill-advised. Such an objection would apply with equal force to the sentence-type analysis outlined above. But

use; they are not simply variants of interrogative forms. Let us assume the question mark indicates a final rise, a common feature of such uses. The force of the objection depends on the validity of several assumptions about sentence type and intonation, including: (i) that all kinds of questioning are *prima facie* incompatible with all degrees of speaker belief; (ii) that rising intonation on a declarative functions simply to confer questionhood; (iii) that declarative 'questions' have the same status as interrogative ones. These assumptions are open to challenge; see Gunlogson (2003, 2008), Bartels (1997) for data and arguments. Without insisting on the adoption of a particular analysis, the point can at least be made that the mere existence of declarative questions does not constitute an objection to the sentence-type approach.

There is much more that could be said about the differences between the two kinds of approaches sketched above, and the merits and demerits of each. In view of the circularity problem, I am inclined toward the sentence-type approach. It is not my goal in this section, however, to argue for the superiority of a sentence-type analysis over a speech-act approach, or vice versa. Instead I want to close by pointing out a larger problem, namely a certain disparity between the requirements for Quality arrived at in this section and the considerations that arose in connection with implicit content in Section 3.1.

In broadening the scope of Quality beyond declaratives used as statements, it becomes necessary to distinguish between types of utterances, as we have seen. Quality becomes more general in the sense that it applies across more utterance types, however those types are defined. It is more specific, though, in the sense that there are multiple clauses of Quality that correspond to the different types. This subdividing on the basis of utterance type leads to some tension between the extension of Quality proposed here with the different sort of generalization discussed in Section 3.1, where it was argued that the domain of Quality should extend to implicitly contributed content. Such implicit content is most naturally conce(in)-203(this) 97(if)-36F103nd e371(e)25dfs6115(371(65(some))-2

4 Quality as a contextual condition

In Section 3.1.2, one rather simple possibility presented itself. Suppose that, instead of viewing Quality as governing only assertive utterances, we implement Quality as a more general rule applying to all (propositional) additions to the context, explicit or implicit. Assuming that implicatures ordinarily constitute additions, the rule will extend to them, as desired. In order to flesh out the implications of such a move, I will start by outlining what I take to be a fairly standard set of assumptions about conditions on assertive contributions in contemporary semantic theory and examine whether and how the proposed extension fits in.

In formal semantics and pragmatics, assertion is routinely modeled in terms of addition to the Common Ground (CG), adapting ideas of Stalnaker (1978). The CG is a set of propositions representing what is mutually taken for granted by the discourse participants. In factual discourses the contents of the CG will correspond to what the participants take as the agreed-upon facts, and will also represent mutual beliefs. An assertion by a speaker is aimed at adding new information to the stock of mutual knowledge. Canonically, utterances of declaratives are taken to function as assertions, and all assertions are assumed to be utterances of declarative sentences, so that, in effect, a speaker uttering a declarative is understood as intending to add new information to the CG.

Typical conditions on assertion are shown in 24. The speaker is expected to be:

- (24)
- a. Informative/non-redundant: offering a proposition that (the speaker believes is) not already in the Common Ground
 - b. Consistent: offering a proposition that (the speaker believes) does not contradict what's already in the Common Ground
 - c. Truthful: offering a proposition that the speaker believes is true and that the speaker has adequate evidence for

24a can be seen as a minimal requirement of Quantity. Quality makes its appearance in 24c, stipulating that the speaker must speak in accordance with his beliefs and evidence.

The three conditions in 24 form a cohesive set from the speaker's perspective, as in Grice's statement of the maxims. All three can be understood as guidelines for the cooperative speaker. There is a difference between 24a–24b and 24c, however, that turns out to be important. 24a–24b make reference to the Common Ground, i.e., to *mutual* assumptions/beliefs, while 24c crucially refers to the speaker's state only.

A standard formulation of requirements 24a–24b in a possible-worlds framework is given in 25:

- (25)
- Let c represent the context set, the set of worlds in which all propositions in the Common Ground are true, and let p be the propositional content of the assertive sentence uttered d , construed as a set of worlds.
 - a. Informativeness requirement: $c \setminus p \neq c$
 - b. Consistency requirement: $c \setminus p \neq \emptyset$

The requirements in 25 are stated in terms of the Common Ground, with the simplifying assumption (often left implicit) that all participants have the same, and accurate, beliefs about the state of the discourse. Given this assumption, the parenthesized references to the speaker's

perspective can be omitted from 24a–24b, which are then defined straightforwardly as CG constraints. But the truthfulness requirement 24c cannot be spelled out so easily; additional structure is needed. Unlike 24a–24b, 24c makes reference to the speaker’s personal belief state, which is not derivable from the CG. Recall that the CG reflects *mutual* beliefs/as.9552 T 36.415 0 867(Recderi)2

beliefs of the interlocutors. Quality should have the result of attributing belief to the contributor, i.e., the speaker, not indiscriminantly to all participants. For one thing, it is possible that not all participants will subscribe to the contributed content. For another, there is a fundamental distinction between being the contributor, the person in the discourse with whom the content originates, and simply accepting the content on someone else's authority. While there may be some sort of Quality principle governing the latter role of accepting, it is not Grice's Quality principle, which pertains to the speaker's knowledge *before* the utterance, not as a result of it.

One necessary step, then, is to disassemble the CG structure into component parts, each representing what one discourse participant is known to assume or believe. If we further stipulate that assumptions correspond to beliefs, then an operation adding a proposition to a participant's set is effectively an addition to that participant's publicly recognized beliefs. Presumably the proposition that the participant in question has the specified belief could uncontroversially pass into the CG, without the content of the belief itself following.

Rather than trying to retrofit the Common Ground with the detail needed to distinguish speaker contributions from addressee acceptances and so on, a revision which takes it some distance from the original conception, we could contemplate starting with a more flexible set of assumptions, returning to Gazdar's proposal for a model. His implementation borrows from Hamblin the notion of a 'discourse commitment' that is not necessarily a belief. A context is a consistent set of propositions; when associated with an individual participant, it is construed as that individual's 'commitment slate', in Hamblin's terminology. Commitments are not explicitly defined, but the important point at the moment is that they are not equivalent to beliefs. The additional step Gazdar takes to spell out the operation of Quality for assertion is to specify that the contributor of an assertion that p is also committed to knowing that p holds, or at least acting that way. Gazdar's informal version of Quality is repeated below:

(27) For any declarative sentence ϕ , assertion of ϕ commits the speaker to $K\phi$.

Adopting this general strategy, the first task would be to generalize beyond the asserted content for attitude attribution. A goal in doing so is to link it to the operation of adding a proposition to a commitment slate, rather than to a specific type of contribution. This amendment has the potential to address the difficulties raised at the end of Section 3.2 involving implicit content added via non-assertive contributions. At the same time, making the attribution explicit allows for its implementation as a default operation, which in turn leaves open the possibility of accounting for the suspension phenomenon by the rather direct method of suspending the default.

I believe the latter approach has much to recommend it, though there are significant issues that arise as well. I leave these to be addressed in future work.

5 Summing up

Summarizing, we have seen that the operation of Quality is distinguished in the following ways:

1. Quality inferences are not cancelable; attempts to deny belief produce a Moore's paradox effect
2. Quality takes precedence; it cannot be traded off against the requirements of other maxims
3. Quality inferences can be globally suspended in non-epistemic contexts, contrasting with the behavior of implicatures in general

4. Quality-like inferences arise for unuttered as well as uttered content, so that attribution of belief to a speaker making a statement will (in the default case) extend to implicatures and presuppositions of the statement
5. Quality can be violated stealthily, without generating violation-related implicatures, unlike other maxim violations (real or apparent), which are inherently public
6. Adherence to Quality must be evaluated with respect to the personal beliefs of individual discourse agents, unlike the other maxims, which can more plausibly be represented as contextual constraints
7. Quality calls for truth-telling and is thus inapplicable on its face to non-assertive utterances, though inferences about speaker attitude arise with such utterances as well

These differences are presented in condensed form in Table 1. I do not wish to claim that the operation of the remaining maxims provides some clear and uniform standard that the behavior of Quality is measured against. However, with respect to the properties discussed, it is fair to group the three together as 'other maxims'.

Property	Quality	Other maxims
Inference cancelable	No	Yes
Inference can lose in maxim clash	No	Yes
Inference suspendable	Yes	No
Inferences about implicit content	Yes	No
Unobtrusive violation possible	Yes	No
Evaluable against CG/context	No	Yes
Crucially refers to truth	Yes	No

Table 1: Summary of differences

This catalogue of distinctive properties provides ample evidence that Quality is not just another, slightly aberrant, conversational maxim but a principle of another sort, deserving of attention in its own right. Together with the points raised in discussion, the list provides a set of desiderata for implementing a more general, 'de-maximized', version of the principle governing matters of quality.

References

- Bartels, C. 1997. Towards a compositional interpretation of english statement and question intonation. Doctoral Dissertation, University of Massachusetts, Amherst.
- Gamut, L. T. F. 1991. *Logic, language, and meaning*. Chicago: University of Chicago Press.
- Gazdar, Gerald. 1979. *Pragmatics : implicature, presupposition and logical form*. New York: Academic Press.
- Gordon, David, and George Lakoff. 1975. Conversational postulates. In *Syntax and semantics*, MLA-IB, 83–106. Los Angeles, CA.
- Grice, H. P. 1989a. Further notes on logic and conversation. In *Studies in the way of words*, 41–57. Cambridge, Mass.: Harvard University Press.
- Grice, H. P. 1989b. Logic and conversation. In *Studies in the way of words*, 22–40. Cambridge, Mass.: Harvard University Press.

