

Is the Symmetry Problem Really a Problem?

Eliza Block

The Symmetry Problem¹ is a simple and powerful challenge to the Gricean explanation of a certain class of Quantity implicatures, to the effect that the Gricean explanation *over-predicts*: for each implicature it correctly generates, it also wrongly generates a conflicting one. Linguists have been aware of the problem for quite some time and have responded by developing accounts of these implicatures that incorporate specialized linguistic competence, and are thus deeply non-Gricean in spirit. [I need a paragraph somewhere about what it means to be ‘Gricean in spirit’.] Recently, there has been an explosion of work on this problem in the semantics literature,² all of which takes the Symmetry Problem seriously as a devastating objection to Grice.

Despite all this attention, the problem remains largely unknown among philosophers. And when confronted with it, philosophers are often reluctant to take it seriously. Discussion of the Symmetry Problem tends to elicit a strong intuition that the problem is illusory. In the first half of this paper, I will investigate this intuition and try to show that it should be set aside. The Symmetry Problem is real.

Despite the fact that no “easy” ways to dissolve the problem pan out, there is one promising avenue that a Gricean could pursue, which is to appeal to the Gricean maxim of Manner to block generation of the unwanted implicatures. Linguists have uniformly dismissed this strategy, on the grounds that there are easy counterexamples that show it won’t work. In the second half of this paper I will argue that the linguists are right

¹This problem has been discussed in passing in the linguistics literature for years (citation?), but it was formalized and named “the Symmetry Problem” by Irene Heim and Kai von Stechow in MIT course lecture notes.

²Fox, Chierchia, Katzir, Spector, Sauerland, Matsumoto, Horn, etc.

to dismiss the Maxim of Manner, but that they are right for the wrong reasons. The arguments they give against solving the problem this way should not convince any defender of the Gricean account. But I will argue that the Maxim of Manner ultimately does not help solve the Symmetry Problem for the Gricean. I will argue for this in a new way, which I hope will be more convincing.

1 The Symmetry Problem

The Symmetry Problem poses a challenge for the traditional Gricean account of inferences like that in (1):

- (1) a. I ate some of the leftovers.
b. $\rightsquigarrow \neg$ [I ate all of the leftovers].

On the Gricean account, the inference in (1) is a conversational implicature generated by the maxim of Quantity. Informally, the implicature arises because the listener reasons (and the speaker expects him to reason) as follows. *The speaker asserted (1a). But there is a relevant stronger (more informative) assertion that she might have made: namely, "I ate all of the leftovers." Since she didn't say this, and we can assume that she is opinionated about the facts and is being cooperative, she must believe this stronger alternative to be false.* It is because of the availability of this reasoning, according to Grice, that (1a) conversationally implicates (1b).

This is a straightforward and intuitive account of the implicature in (1). But it is threatened by the observation that there is another piece of reasoning, exactly parallel to the one just given, that leads to a conflicting (and incorrect) prediction. To see this, note that (1a) partitions logical space into the possibilities in (2).

- (2) a. I ate all of the leftovers.³
b. I ate some but not all of the leftovers.

In order to derive the implicature in (1), we appealed to the fact that (2a) would have been a more informative assertion than (1a). (It's more informative because it rules out the possibility in (2b).) Since the speaker didn't assert (2a), we inferred that she must believe it to be false.⁴ Conclusion: the speaker believes (2b). But now we can see how the Symmetry Problem arises. (2b) would *also* have been a more informative than (1a), since it rules out the possibility in (2a). And the speaker didn't assert (2b), either. So by exactly the same reasoning as above, we ought to infer that the speaker doesn't believe (2b). Conclusion: the speaker believes (2a). In short, the machinery of the Gricean account churns out both (2a) *and* (2b) as predicted implicatures.

Consider another example.

- (3) a. I've met Mary or Sue.
b. \sim I haven't met Mary and Sue.

Just as in the case of (1), the Gricean account of this implicature appeals to the fact that a speaker who asserts (3a) might have said something more informative, namely (4a) below.

- (4) a. I've met Mary and Sue.
b. I've met Mary or Sue but not both.

³This whole debate presupposes, of course, that 'some' is unambiguous and is semantically compatible with 'all'. I discuss the reasons to believe this in a separate paper.

⁴For the sake of simplicity I am intentionally blurring the distinction between *not believing P* and *believing not-P*. This distinction will play an crucial role in any good account of scalar implicatures, but it is not relevant to this discussion of the Symmetry Problem.

Since she didn't, continues the reasoning, she must believe (4a) is false, so she must believe (4b). The Symmetry Problem arises because (4b) is *also* a more informative alternative to (3a). And running the Gricean reasoning on the alternative in (4b) yields exactly the wrong result. If it's false that the speaker has met Mary or Sue but not both, and it's true that she has met Mary or Sue, then she must have met both. So by the very reasoning that allowed us to derive the (genuinely attested) implicature in (3), we also wrongly derive the negation of that implicature.

The Symmetry Problem comes as a surprise, largely because the Gricean explanation of these implicatures *feels* so compelling. If I say I ate some of the leftovers, and you are asked to justify drawing the conclusion that I didn't eat all of them, one very natural thing to say would be "if you had eaten all of them you would have said so." This kind of reasoning is very intuitive; indeed, it seems to be available even to speakers who have never heard of Grice.

To preserve the Gricean account of these implicatures, we need a principled reason why the Gricean mechanism computes an implicature from the alternatives in (2a) and (4a) but not from the alternatives in (2b) and (4b). Throughout my discussion, I'll call alternatives like the ones in (2a) and (4a) the 'good' alternatives: they are good in the sense that if Grice is right, we really do consider them when we compute scalar implicatures. I'll call alternatives like those in (2b) and (4b) the 'bad' alternatives. These are the ones we don't consider, and we need an explanation of why.

Because the Gricean reasoning is so intuitive, it is tempting to think there there must be an easy way out of the Symmetry Problem. The problem depends on the 'good' and 'bad' alternatives being symmetrical in such a way that there is no principled reason to run the Gricean reasoning on the former but not the latter. To dissolve the problem,

then, we would need to show that these pairs are not symmetrical after all; there is some asymmetry that explains why we derive an implicature from only one member of the pair. Importantly, any such asymmetry should be *systematic*—it should apply in all the cases giving rise to the problem—and it should be *available to our general reasoning faculty*. (To give up the latter qualification would be to concede that a Gricean account of scalar implicatures cannot be salvaged.) In what follows, I will try to show that an appropriate asymmetry is harder to find than one might think.

I will consider two different kind of strategies for finding such an asymmetry and thereby dissolving the Symmetry Problem. The first kind of strategy tries to pull an asymmetry out of the very fact that the Gricean reasoning is so intuitively compelling. In section 2 I consider several ways to cash out this strategy, but ultimately argue that it is not helpful to the Gricean cause because it is at worst circular and at best non-explanatory. The second kind of strategy, discussed in section 3, attempts to find an asymmetry in the interaction of the ‘good’ and ‘bad’ alternatives with other Gricean maxims. This is a more promising route: in particular, it seems that the Maxim of Manner might offer a way out of the problem. I argue in the end that that it does not.

2 Intuitive Asymmetries

Could the very fact that the Gricean reasoning is so intuitive provide the asymmetry between the ‘good’ and ‘bad’ alternatives that we need in order to dissolve the symmetry problem? In this section, I’ll consider a few potential asymmetries rooted in the intuitiveness of the ‘good’ alternatives relative to the ‘bad’ ones.

Asymmetry 1: Expressive Choice. Suppose I say “I ate some of the leftovers,” thereby leaving it (semantically) open whether or not I ate all of them. Now, why would you consider “I ate all. . .” as an alternative to my utterance, but not

“I ate some but not all...”? Because suppose I wanted to give you information about whether I finished the leftovers. The only way to convey that I ate all of them would have been to say ‘all’. However, saying ‘some but not all’ is *not* the only way I could convey that I ate only some of the leftovers, because saying ‘some’ also conveys this. In short, there is the following striking asymmetry between ‘all’ and ‘some but not all’: of the three expressions ‘some’, ‘all’, and ‘some but not all’, ‘all’ is the only one that can be used to convey *all*. But ‘some but not all’ and ‘some’ can both be used to convey *some but not all*.

The asymmetry described here certainly exists, but as an attempted resolution of the Symmetry Problem to appeal to it would be viciously circular. The observation that ‘some’ can be used to convey ‘some but not all’—but cannot be used to convey ‘all’—is precisely what we are trying to explain. This asymmetry is due to the implicatures generated by utterances of the word ‘some’. The reason you can convey ‘some but not all’ by saying ‘some’ is that ‘some’ *implicates* ‘some but not all’. The reason you can’t convey ‘all’ by saying ‘some’ is that ‘some’ *does not* implicate ‘all’. Our goal setting out was to explain these facts about implicatures. And we can’t appeal to the phenomenon to be explained in our explanation of it.

Asymmetry 2: Salience. Perhaps the very fact that the Gricean reasoning is so intuitive may itself provide the basis of a solution to the Symmetry Problem. The Gricean reasoning feels intuitive because if someone asserts ‘I ate some of the leftovers’, the alternative ‘I ate all of the leftovers’ *springs readily to mind*, whereas the alternative ‘I ate some but not all of the leftovers’ does not. Independently of any theory of implicature generation, ‘all’ just *is* a salient alternative to ‘some’, whereas ‘some but not all’ just *isn’t*. (Some anecdotal evidence of this difference

in salience: it was only years after Grice published his account of implicatures that anyone even noticed that his reasoning could be run on ‘some but not all’ and thereby generate the wrong predictions.) This asymmetry in salience also generalizes to the other problem cases. If someone says ‘I’ve met Mary or Sue’ we naturally think of ‘I’ve met Mary and Sue’ as a possible stronger assertion; ‘I’ve met Mary or Sue but not both’ comes less readily to mind: ‘A and B’ is a salient alternative to ‘A or B’; ‘A xor B’ is not. Here, then, is a compelling asymmetry between the pairs of alternatives that are supposed to be symmetrical. The ‘good’ alternatives are salient, and the ‘bad’ alternatives are not.

The question of whether salience is a legitimate solution to the symmetry problem will require a more subtle investigation than was required in the previous proposal. There is clearly a very intimate connection between the salience of the ‘good’ alternatives and the fact that they give rise to an implicature. One might worry that this connection is, in fact, *too* intimate. The Gricean is appealing to the maxim of Quantity to explain scalar implicatures: “rule out the stronger thing that could have been asserted.” The symmetry problem consists in the observation that there are two conflicting stronger things that could have been asserted. *Why do we consider only one of them? Because only one is salient.* This sounds suspiciously like a way of repackaging the observation that we do, in fact, consider only one of them. We have not addressed the underlying question: what is the difference between the two alternatives in virtue of which only one is salient and only one is considered by the Gricean machinery?

I’ll return to this matter shortly, because there remains a good deal to say. First, though, I want to consider a third possible asymmetry that builds on the “salience” proposal.

Asymmetry 3: Common Knowledge of Salience. Here's a slightly more indirect version of the above appeal to salience. Whatever the explanation for why the 'good' alternatives are more salient than the 'bad' alternatives, we can explain why we get the implicatures we do by noting that it is common knowledge among members of our linguistic community that these are the salient alternatives. So as a speaker, I can say 'some' with justified confidence that listeners will think of 'all' as an alternative, and won't think of 'some but not all'. Further, I can be confident that my listeners will know that I expect this, and know that I know that they know this. So whatever the initial reason for the asymmetry between 'some' and 'some but not all', all the Gricean needs to do to solve the Symmetry Problem is appeal to the fact that this intuitive asymmetry is common knowledge. Then the Gricean mechanism can exploit it: we get implicatures based on only the 'good' alternatives because those are the alternatives on the basis of which cooperative members of our linguistic community expect one another to reason.

This proposal might appeal to someone who is worried that the "salience" proposal is too close to a mere restatement of the problem. This version drives a wedge between the question of which alternatives are salient and the question of why we get the implicatures we do; it is clearly not just a disguised restatement of the fact to be explained. However, this proposal has its own problems. First, it depends on somewhat implausible empirical facts about what is common knowledge in a linguistic community. The average speaker has presumably never thought about the various alternatives to an utterance of 'some', so this would at best have to be "common knowledge" in a tacit sense. Second, even if the case could be made that the purported common knowledge is actual, appealing to *ad hoc* facts about how members of the linguistic community expect one another to reason

is not a kosher form of defense for the Gricean. Gricean implicatures should be derivable from the maxims of cooperative conversation and the semantic value of the assertion. The above explanation sneaks in facts about how people *usually* reason in response to certain kinds of utterances. To introduce these facts into the equation is tantamount to explaining what implicatures arise by arguing that it is common knowledge what kind of implicatures *usually* arise in response to this kind of assertion, and speakers can be confident that listeners will assume that the speakers intend to convey the same kind of implicature as usual. Which is to explain what implicatures arise now on the basis of what implicatures usually arise. Not great.

So, dispensing with the third proposal, let's return to the "salience" proposal, which is, in my view, by far the most interesting of the different ways to cash out an intuitive asymmetry between the 'good' and 'bad' alternatives. I have raised the concern that the salience of the 'good' alternatives is too intimately connected with the fact that they yield implicatures to be a satisfying *explanation* of this fact. I called it a mere repackaging of the question we set out to answer. But one could easily resist this characterization of salience. After all, there are two phenomena here: (a) the 'good' alternatives are more salient than the 'bad' alternatives, and (b) in computing implicatures we consider the 'good' alternatives and not the 'bad' ones (that is, if the Gricean theory is right). And of these two phenomena, salience may well be explanatorily prior. Perhaps the salience of the 'good' alternatives is due to some deep and mysterious fact about our brain wiring, and to explain it would require some kind of neurological or psychological investigation. Whatever the outcome of that investigation, the fact that the 'good' alternatives are more salient could well explain why we get the implicatures we do. On this way of looking at the matter, the fact that 'some' implicates 'not all' would be explained by a Gricean story about cooperation and general reasoning, modulo

brute facts about our psychology that prevent us from seeing certain alternatives and thus from taking them into consideration when computing implicatures.

The first thing to notice about this kind of theory is that it's already a significant departure from the spirit of Grice's proposal. Grice aims to show how conversational implicatures arise from the meanings of our words and some minimal assumptions about the nature of cooperative behavior. An alien familiar with the semantic value of 'some' and the cooperative maxims should be able to work out the implicature that results from an assertion of 'I ate some of the leftovers'. Once we supplement the theory with dependence on our brain wiring, the alien is out of luck. Further, this answer to the problem carries the uncomfortable suggestion that if we were more perfect reasoners and not subject to physical limitations, we *would* derive implicatures from the 'bad' alternatives; that we simply *neglect* to run the Gricean machinery on alternatives that we (idiosyncratically) fail to notice. Moreover, until we have an explanation (or at least a clear statement) of the psychological facts on which we are relying, we are without a *predictive* theory of implicatures.

A second thing to note is that whatever the explanation for the difference in salience between 'all' and 'some but not all', it is connected to other linguistic phenomena. For example, in each pair of alternatives, the 'good' one (*some, and*) is lexicalized, but the 'bad' one (*some-but-not-all, exclusive-or*) is not. This is true cross-linguistically. (reference?) There are also other related linguistic effects. For example, the "not only X, but Y" construction is only felicitous when Y is the 'good' alternative to X. Compare:

- (5) a. I ate not only some, but all of the leftovers.
b. #I ate not only some, but some but not all of the leftovers.

Given that the distinction between the ‘good’ and ‘bad’ alternatives crops up linguistically in various ways, it seems natural to think the distinction may be part of our linguistic competence. And once we think that, a linguistic solution to the problem of scalar implicatures might look more attractive than a Gricean story supplemented by brute psychological facts.

Those are two further reasons to be skeptical of the salience solution (the latter, admittedly, is rather hand-wavy). But a dedicated pro-Grice salience-ist could perhaps still fail to be dissuaded. So, what happens if we pursue the conversational-maxims-applied-relative-to-brute-psychology route? Our work isn’t done. It’s not enough to say that we only consider the salient alternatives, because we need to say something about the level at which salience applies. How does the psychological mechanism responsible for salience interact with the Gricean reasoning? Does a whole sentence end up being a salient alternative to what was uttered, so we then derive its negation as an implicature? Or does the psychological mechanism present certain *terms* as salient alternatives to other (scalar) terms, then allow us to build up the appropriate sentential alternative by some kind of global replacement strategy?

If the former, then we are appealing to *a lot* of brute facts about salience (brute facts which bear a surprising and unexplained resemblance to one another). Moreover, by taking the sentential route we abandon any hope of a predictive theory. A more parsimonious strategy would be to appeal to the brute psychological fact that ‘all’ is more salient than ‘some but not all’ as an alternative to ‘some’, and ‘and’ is more salient than ‘x-or’ as an alternative to ‘or’, and so on. This is essentially the strategy adopted by Larry Horn (1972), who proposed that scalar implicatures are computed relative to **Horn scales**, of which <some, all> and <or, and> are examples. The Horn scales themselves are hard-coded into the lexicon. From them, we can generate the sentential

alternatives that serve as input to the Gricean machinery.

Even once we've appealed to Horn scales, our solution to the Symmetry Problem is not complete. We still need to explain how the sentential alternatives are generated from these scales. I say "I ate some of the leftovers" and <some, all> is a Horn scale; how, from this, does it follow that "I ate all of the leftovers" is an alternative utterance that gets considered by the Gricean machinery? This turns out to be a harder problem than it appears, and it is one that has been much discussed in the linguistics literature. (See, for example, Sauerland, Chierchia, Fox, Spector.) I won't go into this in detail, but will just mention that the obvious strategy of globally replacing occurrences of 'some' by 'all' runs into trouble when you have embedded scalar terms. For example, consider the following sentence which contains the scalar term 'or' embedded under 'some'.

(6) Some people who smoke or drink develop health problems.

Globally replacing one or more scalar terms in (6) by the stronger members of their scales delivers three alternatives:

- (7) a. Everyone who smokes or drinks develops health problems.
b. Some people who smoke and drink develop health problems.
c. Everyone who smokes and drinks develops health problems.

Would an utterance of (6) typically implicate the negation of all of the sentences in (7)? Answer: it would implicate the negation of (7a), but it clearly would not typically implicate that nobody who smokes *and* drinks develops health problems, which is the negation of (7b). On the other hand, it would arguably implicate that some people who *just* smoke or drink (that is, who smoke or drink but don't do both) develop health

problems. So it seems we get the right result for the embedded scalar when we embed the negation under the quantifier, but not when we negate the whole sentence. But the Maxim of Quantity doesn't seem to allow for negating sub-sentential components.

I mention this example to show that cashing out the 'salience' strategy in terms of Horn scales gets the defender of Grice into uncomfortable territory. There's an enormous amount more to be said about how (and whether) a proposal based on Horn scales will ultimately work, but it looks pretty likely that the details of such an account will involve some appeal to specialized linguistic competence of a kind incompatible with the spirit of Grice's theory. Even if we set aside all of the concerns about the 'salience' strategy that are based on its insufficient *explanatory* power, once we start working out the details of the solution to the point necessary to retain a predictive theory, we're on a slippery slope into non-Gricean territory.

3 Gricean Asymmetries

Perhaps the initial intuitive plausibility of the above attempts to introduce an asymmetry explains some philosophers' resistance to taking the Symmetry Problem seriously. In the end, I don't think any of these suggestions pan out. But there is another important way one might try to resist the Symmetry Problem: by appealing to other Gricean maxims. Here's an outline of how this would go.

According to Grice, 'some' implicates 'not all' because of an interaction between the maxims of Quantity and Quality. When a speaker asserts 'I ate some of the leftovers' he is violating the maxim of Quantity by not providing the relevant information that he ate *all* of the leftovers. Under what circumstances would a cooperative speaker violate this maxim? Answer: when obeying it would force him to violate another maxim, for example the maxim of Quality ("make your contribution true"). In short, we conclude

that the speaker doesn't believe he ate all of the leftovers because we explain his violation of Quantity by appealing to his effort to obey Quality.

Recall the the problem we're trying to solve is to explain why we never derive a parallel implicature from the fact that the speaker didn't assert 'I ate some but not all of the leftovers'. What if asserting this would have violated some *different* Gricean maxim, maxim X? If the 'bad' alternative would have violated maxim X, then there is no need to appeal to Quality to explain why the speaker didn't assert it. So here's an outline of the proposal: Show that the 'bad' alternatives systematically violate some maxim besides Quality, but the 'good' alternatives don't. Then we get the asymmetry we need. Listeners have to appeal to Quality to explain why the speaker didn't assert the 'good' alternative, but they can explain why the speaker didn't assert the 'bad' alternative by appealing to maxim X. Then for all the listener knows, the speaker thinks the 'bad' alternative is true but didn't assert it because he's obeying maxim X. So it would be unreasonable to conclude that the speaker doesn't believe the 'bad' alternative. So no implicature arises.

What could maxim X be?

Maxim of Relation. The Maxim of Relation says "make your contribution relevant."

What if the 'bad' alternatives systematically violated this maxim? If the 'good' alternatives are always relevant and the 'bad' alternatives are always not relevant, then we have an explanation, within the Gricean theory, for why we don't derive the implicature that the 'bad' alternatives are false.

This is a non-starter. I leave it as an exercise for the reader to say why.

Maxim of Manner. Notice that 'some but not all' is awkward, long, bulky, compared to 'all'. Also, 'A or B but not both' is awkward, long, bulky, compared to 'A and

B'. The maxim of Manner includes the clause "make your contribution brief." The 'bad' alternatives seem to be less brief than the 'good' alternatives (granted, based on our sample of only two pairs). But what if the pattern generalizes? What if the 'bad' alternatives systematically violate the brevity condition of the Maxim of Manner? Then we could break the symmetry in exactly the way outlined above. We could say that when a speaker asserts 'I ate some of the leftovers', we don't get a Quality implicature from the alternative 'I ate some but not all of the leftovers' because asserting this would have violated Brevity. On the other hand, asserting 'I ate all of the leftovers' would *not* have violated Brevity, so to explain why the speaker didn't assert *this* we have to appeal to Quality, and we thereby get the implicature that it is false.

Linguists working on the Symmetry Problem have considered appealing to Brevity as a possible way out, but seem to have unanimously rejected this solution. As far as I can tell, the rejection is based on cases where we *do* get a scalar implicature from consideration of a less brief alternative. In other words, it is based on the claim that the pattern noted above does not generalize. In what follows, I discuss these cases and show that they are very weak. They can (and should) easily be explained away by anyone minimally committed to defending the Gricean program. And it is very difficult to find better examples showing that appealing to Brevity won't solve the problem. Furthermore, there is reason to think that this failure is not a coincidence. Roni Katzir (cite) has proposed a promising account of the generation of scalar implicatures that centrally appeals to what he takes to be a formalization of the Brevity clause of the Maxim of Manner.

So where does that leave us? Does Brevity get the Gricean out of the woods? I

will argue that it does not. The Gricean statement of Brevity (“make your contribution brief”) is insufficiently precise to deliver the fine-grained distinctions that would be needed to solve the Symmetry Problem. It does not get us a *predictive* theory of this particular class of implicatures. Katzir’s account fills this void—he spells out the details needed to predict the right scalar implicatures and not the wrong ones. However, Katzir’s account does not vindicate the Gricean, because his account is a fundamentally non-Gricean one. It appeals to details of syntactic structure that are decidedly outside the realm of general intelligence. In short, the Maxim of Manner can only help solve the Symmetry Problem if it is precisified in a way that introduces specialized linguistic competence; the result is no longer a Gricean account.

4 Brevity

Proposal: The symmetry between the ‘good’ and ‘bad’ alternatives can be broken by appeal to the maxim of manner. The ‘bad’ alternatives violate the maxim of manner, so we do not need to suppose that the speaker believes they are false in order to explain why she didn’t assert them. As applied to the two examples we’ve looked at so far, this explanation looks plausible. We’d have to say that “I ate some but not all of the leftovers” violates the maxim of manner, whereas “I ate all of the leftovers” does not; similarly, “I’ve met Mary or Sue but not both” violates the maxim of manner, but “I’ve met Mary and Sue” does not.

What exactly is the maxim of manner doing here? Two possibilities would both get the right results in the above examples. (1) The maxim of manner rules an alternative assertion out of consideration if it is strictly longer than the actual assertion. (2) The maxim of manner rules an alternative assertion out of consideration if it is strictly longer than another alternative assertion. A third possibility would also get the right result

in these cases, but is implausible for other reasons: (3) The maxim of manner rules an alternative assertion out of consideration if it is strictly longer than some fixed length. The third possibility is implausible because assertions of a wide variety of lengths and complexities can give rise to scalar implicatures. So it is far more promising to think that the maxim of manner is indexed to the complexity of the actual assertion. You could see it this way: the actual assertion fixes a level of complexity that the speaker is willing to use. Alternative assertions that are more complex are thus ruled out from consideration. (This makes option (1) look the most plausible as a way of spelling out the contribution of manner. It is hard to see how (2) would be predicted in the Gricean system, although there could be similar systems depending on general reasoning alone that could take advantage of a comparison between the complexity of contrasting alternatives.)

So we have two plausible precisifications of the Brevity response to the Symmetry Problem:

- B1. If A' is a stronger alternative to an assertion A , A' is ruled out from consideration in the computation of Quantity implicatures if A' is strictly more complex than A .
- B2. If A' and A'' are stronger alternatives to an assertion A , A' is ruled out from consideration in the computation of Quantity implicatures if A' is strictly more complex than A'' .

One way to show that the Brevity response fails as a solution to the Symmetry Problem would be to show that there are cases where a Quantity implicature arises based on an alternative A' that violates one or both of the above precisifications. If there were a case where a quantity implicature arises from an alternative A' despite A' being more

complex than A , or than another alternative A'' , that would show that the maxim of manner (so precisified) cannot be the explanation of our failure to derive implicatures from the ‘bad’ alternatives discussed in section 1.

In the linguistics literature on scalar implicatures and the Symmetry Problem, it is almost universally taken for granted that the maxim of manner response does not solve the Symmetry Problem. And the evidence generally taken to support this consists of examples of exactly this kind. Unfortunately, these examples are not convincing. In what follows, I will go through the best cases from the linguistics literature, as well as other cases that seem promising on the surface, and show that none of them really work to show that the maxim of manner response fails. In light of this, the dialectical situation seems to favor the Gricean. Or at least the burden remains on the anti-Gricean to show that the Symmetry Problem survives.

Perhaps the most cited counterexample to the Brevity solution comes from Matsumoto (1995). He points out that an utterance of (8a) would give rise to the implicature in (8b).

- (8) a. It was warm yesterday and it’s going to be a little bit more than warm today.
b. $\rightsquigarrow \neg$ [It was a little bit more than warm yesterday and it’s going to be a little bit more than warm today].

This example seems to show that it is possible to get an implicature from consideration of a more complex alternative. It would thus falsify B1, the first of the above precisifications of Brevity, which says that we don’t consider alternatives that are more complex than what was asserted. If principle B1 is false, then we can’t appeal to it to explain why we don’t get implicatures from the ‘bad’ alternatives, so (unless we resorted to the

less-plausible B2) the Symmetry Problem would stand.

The problem with (8) as a counter-example to the Brevity solution is that it's pretty easy to explain it away. It is natural to think that different levels of brevity/conciseness are appropriate in different contexts, and that a speaker can, in various ways, indicate what standard of brevity she takes to apply in the current conversation. One obvious way to show that a lax standard of brevity applies is to actually use a wordy, awkward, or bulky expression. In Matsumoto's example, the speaker indicates that she considers "a little bit more than warm" acceptable by actually *uttering* it in (8a). So it is consistent with a principle of brevity that we should derive an implicature from considering an alternative that contains that expression. After all, by uttering it, the speaker put it on the table.

At worst, then, Matsumoto's example shows that we would need to modify our Brevity principle(s) to include a provision for alternatives that include expressions that are in use in the conversation. Or we could leave our Brevity principles alone and tinker with what counts as "more complex," treating expressions that are currently in use or otherwise marked as acceptable as more simple than their grammatical structure might suggest. (As I'll discuss, Roni Katzir includes such a provision in his structural alternatives account; he builds special permission for "in-use" expressions into his definition of complexity.) Either way, the Gricean has a straightforward fix for examples like Matsumoto's, and it isn't even a bothersomely ad-hoc fix. So far, then, the Brevity solution stands.

Here's another counterexample to the Brevity solution, one that threatens both B1 and B2. (I believe this example is due to Fox, but I don't have the reference handy.) It depends on the somewhat controversial assumption that number words like 'three' are scalar terms. On this view, 'three' asymmetrically entails 'two', and is compatible with

‘four’. So three is compatible with both *more than three* and *exactly three*, each of which is more informative than *three* because each rules out the other. The counterexample comes from the observation that ‘more than three’ is the ‘good’ alternative, and ‘exactly three’ is the ‘bad’ alternative. For example, we easily derive the implicature in (9):

- (9) a. I have three children.
b. $\rightsquigarrow \neg$ [I have more than three children].

But we don’t derive the implicature that \neg [I have exactly three children]. This counterexample seems to falsify B1 because we’re deriving an implicature from a more complex alternative. It also poses trouble for B2 because ‘more than three’ looks more complex than ‘exactly three’. (Of course that depends on our metric of complexity, which we’ve said nothing about. Here I’m just counting words for the sake of argument. Obviously that would not be a very satisfying final theory.)

The dialectical situation here is a little convoluted, so let me summarize. The pro-Gricean has proposed something like B1 or B2 as a principle that, if correct, gives us an asymmetry that could solve the Symmetry Problem. The anti-Gricean has put forward (9) to show that B1/B2 is not a correct principle: implicatures arise in violation of it. (9) is a little shaky because it rests on some significant assumptions (most importantly that numbers are scalar terms, which many people deny). Now, I’m trying to argue that it’s hard to find real counterexamples to the Brevity principles. So I’m going to grant all the assumptions that are needed to solidify (9) as a counterexample. Let’s suppose that numbers are scalar, and (to challenge B2) that ‘more than three’ is more complex than ‘exactly three’. *Even granting all that*, I will argue, (9) does not refute the Brevity solution.

Here’s why. In order for (9) to be a counterexample in the first place, we had to grant that number terms are scalar, so ‘three’ means something like ‘three or more’. (If we didn’t assume this, we would have to treat the inference in (9) as an entailment, not an implicature. But for the sake of the counterexample we need it to be an implicature.) Okay, so ‘three’ means ‘three or more’ and ‘four’ means ‘four or more’, which is logically equivalent to ‘more than three’ (provided we’re only thinking about integers, which is, I hope, a safe assumption given that we’re counting children). But that means that there’s a way to characterize the alternatives in this example that *doesn’t* violate either version of our Brevity principle: the ‘good’ alternative to ‘three’ is ‘four’, and the bad alternative is ‘exactly three’. Instead of the implicature in (9), we get the one in (10):

- (10) a. I have three children.
 b. $\rightsquigarrow \neg$ [I have four children].

On the assumption that ‘four’ is scalar, (10b) is equivalent to (9b). And since ‘four’ is no more complex than ‘three’ and is less complex than ‘exactly three’, the apparent conflict with the Brevity principles dissolves.

Another possible counterexample to the Brevity response:

- (11) a. I’ve met *Mary*.
 b. $\rightsquigarrow \neg$ [I’ve met Mary and Sue].

(To get this implicature you have to imagine this conversation happening in a context where Mary and Sue are both salient people that the speaker might have met. Imagine that (11a) is said in response to the question “Which of my sisters have you met?” Also, note that ‘Mary’ is focused.) Again, this looks like an implicature that requires

consideration of a more complex alternative, in conflict with at least B1 above.

I won't go into this in much detail, but I think this example fails to be convincing as well. It seems to be an inference that is a result of the focus on 'Mary'. The most common view about how this works is that the context sets up a set of focus alternatives, which in this example would be the set of the speaker's sisters. Then the assertion of (11a) implicates (actually, is this usually considered to be an implicature or a semantic effect? I'm not certain) the negation of the alternatives created by substituting a member of the contrast set. So we'd get an inference that the speaker has not met *Sue*, and an inference that he has not met *Nancy*, and so on. It's not clear to me what relation these focus inferences bear to scalar implicatures (they seem very similar), but it is natural to think that the correct specification of the alternatives does not involve the conjunction, so makes each alternative equally simple as the original.

The upshot is that it's harder than one might hope to find clear counterexamples to the Brevity solution to the Symmetry Problem. However, it should also be clear at this point that the Brevity solution is not sufficiently fleshed out to yield real predictions. Most importantly, we have not said what metric of complexity we are using to determine the relative brevity of different alternatives. But perhaps this gap can be filled by a theory proposed by Roni Katzir in his paper "Structurally Defined Alternatives." If Katzir is right, there is a straightforward syntactically defined metric of complexity that yields the right predictions about scalar implicatures.

5 Structural Complexity

[This section is rather abbreviated in this draft; I intend to flesh it out in a later version.]

Katzir proposes using the following metric of complexity: S' is at least as brief as S iff S' can be derived from S by a series of pruning and replacement operations on S 's syntactic

tree. Very roughly, any node in the syntactic tree is allowed to be replaced by a member of the substitution source, which includes lexical items and “in-use” expressions (such as ‘a little bit more than warm’ in the Matsumoto example in (8)). For example, in “I saw Bill and John” we can replace the node that contributes “Bill and John” by a new node containing “Mary,” resulting in “I saw Mary.” The latter sentence is strictly simpler than the former, because the rules allow for a transition in that direction but not in reverse. (The reverse transition is blocked because “Bill and John” is not a member of the substitution source.)

The theory, then, is that scalar implicatures arise from the Maxim of Quantity applied relative to sentential alternatives that are at least as brief, according to this metric, as what was asserted. So it’s B1 from above combined with a structural definition of brevity. Now we have a true predictive theory, and its predictions look pretty good. For example, it successfully predicts the implicature in (12):

- (12) a. Everyone who likes John but not Mary is an idiot.
 b. $\rightsquigarrow \neg$ [Everyone who likes John is an idiot].

The alternative negated in (12b) is strictly briefer than the assertion in (12a), because it can be generated from (12a) by replacing the node that contributes “John but not Mary” with a node containing just “John.” This is particularly nice, because a lot of other theories of scalar alternatives fail to make this prediction. An account that appeals to Horn scales, for example, would have trouble making this prediction because $\langle \text{John, John but not Mary} \rangle$ would not usually be considered a scale.

As nice as this theory looks, however, I am ultimately skeptical that it can come to the rescue of the pure Gricean. The reason is that by appealing to this kind of metric

of complexity, we have departed rather sharply from the realm of general reasoning. It looks like the kind of computation required to generate the alternatives that must serve as input to the Gricean machinery crucially involves a syntactic component. So at best, we end up with a Grice-like theory that appeals centrally to specialized linguistic competence in a crucial stage of the reasoning. This might be the correct theory, but it is far from what Grice had in mind.

6 Conclusion

Is the Symmetry Problem really a problem? Yes. It is a real problem for the traditional Gricean picture and its account of scalar implicatures. There seems to be no way to solve it while preserving predictiveness and shunning reliance on specialized linguistic competence.