

Forthcoming in François Recanati, Isidora Stojanovic, Neftali Villanueva, eds., *Context-dependence, Perspective and Relativity in Language and Thought*, Mouton.

## **Indirect Discourse, Relativism, and Contexts that Point to Other Contexts**

Christopher Gauker

Department of Philosophy

University of Cincinnati

*Abstract:* Some expressions, such as “all” and “might”, must be interpreted differently, relative to a single context, when embedded under “says that” than when unembedded. Egan, Hawthorne and Weatherson have appealed to that fact to argue that utterance-truth is relative to point of evaluation. This paper shows that the phenomena do not warrant this relativistic response. Instead, contexts may be defined as entities that assign other contexts to contextually relevant people, and context-relative truth conditions for indirect discourse sentences can be satisfactorily formulated in terms of such contexts.

Indirect discourse is different from direct quotation. In an indirect discourse sentence, the sentence that follows the complementizer “that” need not be the very sentence that the

---

This research was supported by grants from the Taft Research Center of the University of Cincinnati. A discussion with Michael Glanzberg in Oslo provoked me to get clearer about some of the matters discussed in this paper, and I am grateful to him for it.

speaker spoke. For instance, the sentence, “Silas said that everyone at the party was wearing a funny hat” might be true even if the sentence Silas actually uttered was “Everyone is wearing a funny hat”, not “Everyone *at the party was* wearing a funny hat.” A semantic theory of indirect discourse sentences ought to have something to say about how the wording in a true indirect discourse sentence may vary from the wording of the speech act that it reports.

In this paper I have four main tasks. First, I will illustrate the way in which inattention to this issue can produce fundamental confusions in semantics. I will do this by criticizing an attempt to defend a kind of relativism on the basis of claims about indirect discourse. Second, I will explain how a semantics for indirect discourse sentences may be formulated in terms of contexts in such a way that a context may *point* to other contexts. The context relative to which we evaluate an indirect discourse sentence will *point* to the context relative to which we evaluate the utterance used in the act of speech that the indirect discourse sentence reports. Third, I will highlight some of the assumptions about the nature of semantics and the nature of contexts that I will have made along the way. Fourth, in an appendix, I will sketch a precise semantics for a simple language that incorporates the semantics I will have described informally in the second part. A couple of related topics will be touched on along the way as well.

### **1. An unpersuasive argument for relativism**

The right way to accommodate context-relativity in semantics, I will assume, is to formulate one’s semantic theory as a recursive definition of the conditions under which sentences are true *relative to* a context. Thus, sentence truth is a *relation* defined

recursively on an inductively defined set of pairs consisting of one sentence and one context.

*Utterances* of sentences, however, may be true or false simpliciter. That is, truth for utterances remains a *property* and is not a relation to a context. The relation between truth-in-a-context for sentences and truth simpliciter for utterances may be stated as follows:

An utterance of a sentence is true (simpliciter) if and only if the sentence uttered is true in the context that pertains to the utterance.

In view of this, one thing we will have to explain, if not as part of our *semantic* theory proper, then as part of our larger theory of *language*, will be what it takes for a given context to be *the* context that *pertains* to an utterance. (Instead of saying that a context pertains to an *utterance*, we can say that it pertains to the *conversation* in which the utterance occurs. We may individuate conversations in such a way that a context that pertains to an utterance pertains as well to every utterance that occurs in the same conversation.)

If this is the right way to look at things, then it is important to distinguish between *contexts* and *situations*. A *context* is a formal structure of some kind containing a value for each of a variety of parameters on which the truth value of a sentence may depend. A *situation* is a course of events and arrangement of objects containing one or more utterances of sentences. (Or at least, the only situations that we will care about will be those containing an utterance.) What we will look to in trying to decide which context is the one that pertains to a given utterance will be features of the situation in which it occurs.

One kind of *relativism* in semantics may be defined as the thesis that utterances too are not just true or false simpliciter but only relative to an evaluator. There are passages in Egan, Hawthorne and Weatherson 2005 (pp. 131, 154; abbreviated EHW) and in Egan 2007 (p. 2) that seem to define relativism in just that way. The truth of an utterance, it is said, is relative to a *context of evaluation*. This does not have to be the kind of relativism that the friends of science will abjure, for the context of evaluation might make a difference only for certain kinds of sentence, and those kinds may exclude the verdicts of science. which we expect to be true in all contexts of evaluation if they are true in any.

Still, anyone who expects there to be a precise, recursive semantics for natural language must regard this “definition” of relativism as at best a derivative characterization, and not a definition. I take for granted that the set of utterances does not form an inductive set. Nor does it form an inductive set up to a maximum length or complexity of utterance. (Utterances cannot suitably be analyzed into uttered components. The negation of a sentence may be uttered though the sentence negated has never been uttered.) So it does not make sense to propose a recursive definition of truth relative to context of evaluation for utterances.

Rather, the relativist thesis should be defined as the claim that the contexts relative to which we evaluate *sentences* must include a parameter for point of evaluation. And then the relativist can maintain that what is special about this parameter is that no single point of evaluation uniquely *pertains* to a given utterance. So the semantic value of an *utterance* is irreducibly a relation between an utterance and a point of evaluation. The approach along these lines favored by EHW and Egan is to define (in Egan’s

terminology) *centered-worlds propositions* as sets of triples, called *centered worlds*,  $\langle w, t, i \rangle$ , consisting of a world  $w$ , a time  $t$ , and an agent  $i$  (EHW, p. 157; Egan p. 5).

Although EHW and Egan do not do so, one could write a recursive semantics as a theory of the conditions under which a sentence expresses a centered-worlds proposition.<sup>1</sup> We can suppose that an utterance is true relative to an agent-time pair  $\langle t, i \rangle$  if and only if the centered-worlds proposition that the sentence uttered expresses contains  $\langle @, t, i \rangle$  ( $@$  being the actual world). So if we identify contexts of evaluation with agent-time pairs, the first characterization of relativism, according to which utterance truth is relative to context of evaluation, follows from the second, according to which sentences express centered-world propositions.

Some popular arguments for a relativist semantics turn on matters of indirect discourse. In particular, EHW and Egan tell “eavesdropping” stories such as the following: Moriarty, in London, says to his assistant,

(1) Holmes might be in Paris.

Assuming that neither Moriarty nor his assistant knows that Holmes is not in Paris, and assuming that Moriarty has some reason to think that Holmes might be in Paris, we may

---

<sup>1</sup> EHW want to say that a sentence expresses the same “proposition” in every context but that “the truth-value of that proposition is contextually variable” (2005, p. 154). But the varieties of context-relativity they actually discuss are very limited – only those that they want to redesign as relativity to points of evaluation. For instance, they do not discuss demonstratives, or the context-relativity of quantifiers, or standards for comparative adjectives. Perhaps they think that all of these are determined by the choice of world, time and agent. Alternatively, they might be prepared to build additional parameters into the  $n$ -tuples that propositions are supposed to be sets of, or they might be prepared to say that in other respects even the expression of centered-worlds propositions may be relative to a context.

suppose that Moriarty's utterance of (1) is true. Meanwhile, Holmes and Watson are listening in on the other side of the wall. Watson says to Holmes,

(2) Moriarty says that you might be in Paris.

Likewise, we may suppose that Watson's utterance of (2) is true. (EHW have "believes" in place of "says" in their version of (2), p. 155.)

The problem is that on a certain plausible understanding of "might" and a certain understanding of indirect discourse, (2) seems to be false. We may assume that "might"-sentences such as (1) are to be evaluated relative to contexts that specify domains of possibilities of some kind.<sup>2</sup> (1), then, is true in the context that pertains to Moriarty's utterance of it, since there is a possibility in the set of possibilities specified by that context such that Holmes is in Paris in that possibility. But the context that pertains to Moriarty's utterance of (1) is different from the context that pertains to Watson's utterance of (2). Since Holmes is right there in front of Watson in London, the set of

---

<sup>2</sup> EHW and Egan call the kind of "might"-sentence they are interested in *epistemic modals*. I myself am not so sure that there is a special class of sentences or modal operators that deserve to be called "epistemic". Rather, it may be that every sentence of the form "It might be that *p*" is true or false only relative to a contextually determined set of possibilities. In some conversations, the context that pertains to that conversation determines a set of possibilities comprising all and only those that are compatible with what the speaker knows or with what any member of a certain set of people, including perhaps the speaker, knows (so that a possibility is not included if there is someone in the set such that it is incompatible with what he or she knows). But on other occasions, the context pertinent to a conversation will determine a set of possibilities in a different way. For instance, it might be the set of possibilities compatible with the interlocutors' carrying out some plan that they have made ("We might wait until next week to get started") or the set of possibilities compatible with their conforming to some rules they have laid down ("We might vote on this before the committee reports").

possibilities pertinent to Watson's utterance presumably does not contain a possibility in which Holmes is in Paris. And likewise, we should not say that Moriarty *said* that there was such a possibility in that domain of possibilities. But if we take the domain of possibilities relative to which we evaluate the "that"-clause in (2) to be those that are pertinent to Watson, then, apparently, that is what we will interpret (2) as saying Moriarty said. So (2) will be false.

As relativists, EHW and Egan propose to avoid this result by taking the proposition a sentence expresses to be a centered-worlds proposition. Since the proposition that is expressed by the sentence that follows the "that"-clause contains not only triples  $\langle w, t, i \rangle$  in which the center  $i$  is the speaker but also triples  $\langle w, t, i \rangle$  in which  $i$  is the person to whom an act of saying is attributed, there is supposed to be the possibility of defining the saying-relation in such a way that what Watson's utterance of (2) means is that Moriarty said that for at least one of the possibilities pertinent to *Moriarty* (not Watson) Holmes is in Paris (EHW 2005, p. 158; Egan 2007, p. 9). The basic idea is that (somehow) the relativist can allow that the utterance of (2) is evaluated from the point of view of the speaker of (2) by deciding how the attributee, the subject of (2), Moriarty, evaluated the utterance of (1) from *his* point of view.

Actually, I am not at all confident that what EHW and Egan say even makes good sense. While I am focusing on "says", EHW and Egan focus on "believes". What EHW *say* about belief sentences is this:

When one says that *a believes that b might be F*, one says that *a* believes the proposition *b might be F*. And *a* believes that proposition iff *a* believes it is consistent with what they know that *b* is *F*. (2005, p. 158)

And what Egan *says* is this:

When I believe that it might be that  $P$ , I believe something about my particular situation relative to the evidence—namely, that none of the evidence that's within my epistemic reach rules out  $P$ . (2007, p. 9).

But as far as I can see, EHW and Egan have given us no clear basis for these conclusions. Since Egan's solo paper is more explicit in its analysis of "might"-sentences, I will focus on what he says there.

Apparently, Egan thinks the account of belief expressed in the above quotations follows from three premises:

- (i) His account of "might": "It might be the case that  $p$ " is true relative to a centered world  $\langle w, t, i \rangle$  iff it's compatible with everything that's within  $i$ 's *epistemic reach* at  $t$  in  $w$  that  $p$  (p. 8).
- (ii) His relativistic conception of propositions: Propositions are sets of centered worlds (p. 5).
- (iii) The assumption that belief (and, I suppose, saying) is a relation between a subject and a proposition (p. 6).

From these premises it is apparently supposed to follow that  $S$  believes (says) that *it might be that  $p$*  if and only if  $S$  believes (says) that it is compatible with everything within his or her epistemic reach that  $p$ .

In particular, Egan needs it to be the case that (a), below, is sufficient for (b):

- (a)  $S$  believes (says) that it is compatible with everything within his or her epistemic reach *that  $p$* .
- (b)  $S$  believes (says) that *it might be that  $p$* .

If (a) is sufficient for (b), then, as desired, it will be sufficient for the truth of (2) that what Moriarty says is, in effect, that there is a possibility within his epistemic reach in



which Holmes is in Paris. Whereas, if (a) is not sufficient for (b), then we will still have no explanation of why (2) should not be deemed false on the grounds that Moriarty did not say that it is compatible with everything within *Watson's* epistemic reach that Holmes is in Paris.

But on the contrary, from (i), (ii) and (iii) it does *not* follow that (a) is sufficient for (b). If (i), (ii) and (iii) hold, then what (b) says is much stronger than what (a) says. What *S* believes (says) according to (a) concerns only the epistemic reach of *S*. But by (i), (ii) and (iii), *S*, in believing (saying) that it might be that *p*, stands in the belief (saying) relation to a certain set of centered worlds, namely, the set of centered worlds  $\langle w, t, i \rangle$  such that it is compatible with everything in the epistemic reach of *i* at *t* in *w* that *p*. So if (a) is sufficient for (b), then the fact that *S* believes that it is compatible with everything in his or her epistemic reach that *p*, as (a) says, tells us, for every *w*, *t* and *i*, whether or not  $\langle w, t, i \rangle$  belongs to a set of centered worlds to which *S* stands in the belief relation. But how can that be? The proposition that it is compatible with everything in the epistemic reach of *S* that *p* – even if construed as a centered-worlds proposition – does not tell us anything about whether it is compatible with everything in the epistemic reach of *anyone else* that *p*. So *what S* believes according to (a) does not imply *what S* believes according to (b), given (i), (ii) and (iii). Since *what S* believes according to (a) does not imply *what S* believes according to (b), surely *S's believing* what *S* believes according to (a) also does not imply *S's believing* what *S* believes according to (b) either.

EHW are driven to their relativism in part by their belief that what they call contextualism cannot account for the phenomena. What they call contextualism is the thesis that the proposition expressed by a sentence is relative to the context in which it is

uttered. The argument against contextualism is reduced to its essentials in Weatherson 2008. In that paper, Weatherson makes very explicit his assumption that if contextualism is true, then context-relative expressions such as “might” must be interpreted in the same way whether or not they are embedded in “that”-clauses. (In his argument on p. 535, this is the conjunction of premises 1 and 2, and in his argument on p. 537, this is the conjunction of premises 4 and 5.) Although the contextualism that I defend is not quite what EHW and Weatherson define contextualism to be (I will not formulate my account in terms of any kind of proposition), I will nonetheless show that Weatherson’s basic assumption is false. Even when “Moriarty says that you might be in Paris” and “You might be in Paris” are evaluated with respect to a single context, it is not the case that the embedded “might” and the unembedded “might” have to be interpreted in the same way.

Nonetheless, I think that EHW, Egan and Weatherson have stumbled upon a semantic question of real interest, and I agree that it may motivate significant revisions to our conception of the semantics of natural language. Very roughly, the problem is that in some respects the “that”-clause in an indirect discourse sentence reflects the perspective of the speaker of the indirect discourse (the attributor), and in some respects it may reflect the perspective of the person to whom saying is attributed (the addressee). What the relativists are picking up on and trying to accommodate are the respects in which the “that”-clause reflects the perspective of the addressee, not the attributor. But as we will see, it is not necessary to resort to relativism to accommodate this distinction.

## 2. Universal quantifiers in indirect discourse

In this paper, I do not want to try to give a semantics for “might”, because that task raises all kinds of additional issues that I do not want to get into here. So let me shift to relatively straightforward quantifiers such as “every” and “everyone”. Like “might” in indirect discourse, “every” in indirect discourse can reflect the perspective of the attributee, as opposed to that of the attributor. But moreover, the attributor can add clarificatory modifiers to the quantifier, and this fact too needs to be accommodated in our semantics.

Suppose that Silas is at the birthday party of his friend Jonas. Silas’s sister Emma is at a different birthday party for a friend of her own. The party for Emma’s friend is a bit listless; so Emma calls Silas on her cell phone to see how things are going where he is. Silas is having a great time and tells Emma excitedly, “There’s a clown here; he’s making animal shapes with balloons!” Emma turns to her friend Dottie, who is also getting a bit bored and says,

(3) Silas says that there’s a clown there.

What Emma does not say to Dottie is,

(4) \*Silas says that there’s a clown here.

So the words that follow “that” in Emma’s utterance are not quite the words that Silas uttered. Emma substitutes “there” for Silas’s own “here”.

The next thing Silas says is “Everybody is wearing a funny hat!” Emma turns to Dottie and says,

(5) Silas says that everybody is wearing a funny hat.

So the words that follow “that” in Emma’s utterance are the very words that Silas spoke. However, Emma might have reported what Silas said differently. She might have said,

(6) Silas says that everybody at Jonas’s party is wearing a funny hat.

She *might have* said (6), but if Dottie knows that Silas is at Jonas’s party, it is not necessary for her to specify in this way that the people Silas is talking about in saying “everybody” are the people at Jonas’s party. Even if Dottie does not know where Silas is, it would not really be wrong for Emma to utter (5) to Dottie, just less than optimally informative.

What the example suggests is that *here/there* and *everybody* are different sorts of cases. In Emma’s report to Dottie about Silas’s saying about the clown, she has to substitute “there” for Silas’s “here”. In other words, the sentence that follows “that” has to be altered to reflect her own point of view, or the context pertinent to her conversation with Dottie. But in Emma’s report to Dottie about Silas’s saying about the funny hats, Emma does not need to substitute anything for Silas’s “everybody” to reflect her own point of view, although she may do so for greater clarity. So it seems we can say at least this: There are two kinds of context-relativities, those that require adjustment to the speaker’s context in indirect discourse (“here” being adjusted to “there”) and those that do not require adjustment to the speaker’s context (“everybody” remaining as it is).

That we need to say something a little more complex than this becomes apparent when we consider *how* a speaker may modify the quantifier in case she chooses to modify it for clarity. One thing the speaker may do is “precisely subclassify the universal domain.” What I mean by that is that if the attributee’s actual words are of the form, “Every *F* is *G*” and the set of *FH*-things *in the whole world* equals the set of *FH*-things in

the domain of discourse pertinent to the attributor's utterance and also equals the set of *F*-things in domain pertinent to the attributee's utterance, then the attributor may say, "*S* said that every *FH* is *G*". That is what Emma would do if she uttered (6) instead of (5) (on the assumption that it is clear what the determiner phrase "Jonas's party" refers to). For the set consisting of anyone in the universe who is attending Jonas's party equals (we may suppose) the set of people in the domain pertinent to Emma's utterance who are at Jonas's party, and also equals the domain pertinent to Silas's utterance, which consists of the people at Jonas's party.

But a speaker does not have to choose between merely quoting the speaker's quantifier (adding no additional modifiers) and precisely subclassifying the universal domain in this sense. The material that the attributor adds to the words of the attributee may serve to clarify without precisely subclassifying the universal domain. Suppose that Julian is showing Ingrid his pottery collection. With a sweep of his hand, he says, "Everything is from Korea." In reporting what he said, we may say,

(7) Julian said that every piece of pottery was from Korea.

Thus, we add the words "piece of pottery" to his "every". But the phrase "every piece of pottery" does not precisely subclassify the universal domain, because the set of pieces of pottery (in the whole world) is not equal to the set of pieces of pottery in the domain pertinent to Julian's utterance, which is confined to the pottery in his collection. Or suppose that Mary is the chair of a committee at her workplace. At the start of the meeting, she looks around and declares, "Everyone is present." Later, one of Mary's co-worker's, may correctly assert,

(8) Mary said that everyone working on a graphic interface was present.

That is something the co-worker can say if the committee consists of the people at her workplace who are working on a graphic interface. The phrase “everyone working on a graphic interface” does not precisely subclassify the universal domain, because the set of people (in the world) working on a graphic interface (for a project somewhere) is not equal to the members of the domain of discourse pertinent to Mary’s utterance.

Thus, the question arises, what determines what material can be added to a universal quantifier in an indirect discourse report on what another person as said? In the next section I will describe, informally, a semantics for indirect discourse that answers this question.

### **3. A semantics for indirect discourse, informally**

For purposes of formulating a semantics of indirect discourse, we may conceive of a context as an entity that assigns other contexts to various pertinent people. In other words, contexts are identified in part by what other contexts they *point* to. (For a precise formulation of this and other definitions in this section, see the Appendix.) So a sentence of the form “*S* said that *p*” will be true in a context *c* *only if* *c* assigns a context *c*(*S*) to *S*. (This is only a weak necessary condition.) Call *c*(*S*) the *context determined for S by context c*. (Here I am using the variable “*S*” as a variable ranging over people’s names. However, for simplicity, I will frequently refer to people and their names interchangeably.)

Also, we will want to suppose that for each of the people to whom a context *c* assigns a context, the context *c* determines a domain of sentences such that relative to *c* that person is considered to have uttered those sentences. Call this the *utterance domain*

for  $S$  in  $c$ . Moreover, we will want to define a relation between sentences such that “ $S$  said that  $p$ ” will be true in a context  $c$  only if there is a sentence in the utterance domain for  $S$  in  $c$  such that that relation holds in  $c$  between  $p$  and the sentence uttered. Whether that relation holds in  $c$  will depend on the content of  $c$  and, in particular, will depend on the content of the context determined for  $S$  by  $c$ . Let us call this relation, still to be defined, *elevation to  $c$  from  $c(S)$* . So we may say that one sentence is an *elevation* of another sentence *to  $c$  from  $c(S)$* .

As I said at the start of section 1, our theory of language will include an account of the conditions under which a context *pertains* to a given utterance. We may impose as a condition that a context must meet in order to count as the context that pertains to an utterance that for each agent  $S$  to whom the context assigns an utterance domain, the sentences in the utterance domain assigned to  $S$  must be sentences that  $S$  has actually uttered. A further condition on  $c$ 's pertaining will be that the context determined for a speaker  $S$  by  $c$  must be the context that pertains to  $S$ 's utterances of the sentences in the utterance domain for  $S$ . (For simplicity, then, I am assuming that only a single context pertains to all of those utterances. In a fuller account, we might wish to allow that a speaker is associated with various sets of utterances, each associated with its own context.)

Putting these devices together, we may say that “ $S$  says that  $p$ ” is true in a context  $c$  if and only if there is a sentence  $q$  such that  $q$  belongs to the utterance domain for  $S$  in  $c$  and  $p$  is an *elevation* of  $q$  to  $c$  from  $S(c)$ . In view of the stated conditions on pertaining, this implies that an utterance of “ $S$  says that  $p$ ” will be *true (simpliciter)* only if  $S$  uttered a sentence  $q$  such that  $p$  is an elevation of  $q$  to the context pertaining to that utterance of

“ $S$  says that  $p$ ” from the context that pertains to  $S$ ’s utterance of  $q$ . (This necessary condition is not also sufficient, because it may happen that not every one of  $S$ ’s utterances is an utterance of a sentence in the utterance domain that  $c$  assigns to  $S$ .)

In these terms, here is how we can account for the fact that sentence (3), “Silas says that there’s a clown there”, is true in the context pertinent to Emma’s utterance of it (Emma’s context), and thus for the fact that Emma’s utterance of (3) is true (*simpliciter*). First, we may suppose that Silas’s utterance of “There’s a clown here” belongs to the utterance domain for Silas in Emma’s context, since Silas did utter that sentence. Second, we may suppose that the sentence “There’s a clown there” is an elevation of “There’s a clown here” to Emma’s context from the context that Emma’s context assigns to Silas. The basis for this assumption will be the fact that the place that “there” refers to in Emma’s context is the place that “here” refers to in the context that Emma’s context assigns to Silas. So the sentence (3), “Silas said that there’s a clown there” is true in Emma’s context, because the sentence “There’s a clown here” is in the utterance domain for Silas in that context and “There’s a clown there” is an elevation of the sentence “There’s a clown here” to Emma’s context from the context that Emma’s context assigns to Silas. Since Emma’s context is the context that pertains to her utterance of (3) and (3) is true in that context, her utterance of (3) is true.

Similarly, but even more simply, we can account for the fact that (5), “Silas says that everybody is wearing a funny hat”, is true in the context that pertains to Emma’s utterance. “Silas says that everybody is wearing a funny hat” is true in Emma’s context, because “Everybody is wearing a funny hat” belongs to the utterance domain for Silas in that context and “Everybody is wearing a funny hat” is an elevation of *itself* to Emma’s



context from the context that Emma's context assigns to Silas. Likewise, (6), "Silas says that everybody at Jonas's party is wearing a funny hat", is true in the context that pertains to Emma's utterance, because "Everybody is wearing a funny hat" is in the utterance domain for Silas in that context and "Everybody at Jonas's party is wearing a funny hat", we may suppose, is an elevation of "Everybody is wearing a funny hat" to Emma's context from the context that Emma's context assigns to Silas. Both (5) and (6) can be true in the context that pertains to Emma's utterance, because "Everyone is wearing a funny hat" is in the utterance domain that Emma's context assigns to Silas, and there are two different sentences that are both elevations of that sentence to Emma's context from the context that Emma's context assigns to Silas.

Often, whether one sentence is an elevation of another to  $c$  from  $c(S)$  will depend both on the contents of  $c(S)$  and on other elements of  $c$ . For example, the reason why "There's a clown there" is an elevation of "There's a clown here" to Emma's context from the context that Emma's context assigns to Silas is that there is a single location  $a$ , such that Emma's context assigns  $a$  to "there" and the context that Emma's context assigns to Silas assigns  $a$  to "here". Likewise, to explain the qualifications that may be added to a quantifier in indirect discourse, as in (6), (7) and (8), we need to define the elevation relation in a way that makes reference to both the context pertinent to the attributor's utterance and the context that that context assigns to the attributee, as I will presently explain.

Toward defining the pertinent principle of elevation, let us stipulate that every context contains a *domain of discourse*. I will suppose that, strictly speaking, the members of domains of discourse are singular terms, not objects such as terms might be

thought to denote. (Think of them as demonstratives with subscripts: *that*<sub>32</sub>.) But sometimes I will write as if domains of discourse were sets of nonlinguistic objects. So for each context  $c$ , there is a domain of discourse  $D_c$ , which is a nonempty set of singular terms. Moreover, an atomic sentence “ $t$  is  $F$ ” is true or false in a context only if  $t$  belongs to the domain of discourse for that context. (It is neither true nor false if  $t$  is not a member of the domain. Our semantics is gappy, or three-valued. However, it is not in general true that if a term occurs in a sentence that is true in the context then the term is in the domain for the context. If “ $t_1$  is  $F$ ” is true in a context, then “ $t_1$  is  $F$  or  $t_2$  is  $F$ ” might be true in that context whether or not  $t_2$  is in the domain for the context.)

I contend that the following principle of elevation holds:

“Every  $FH$  is  $G$ ” is an elevation of “Every  $F$  is  $G$ ” to  $c$  from  $c(S)$  if and only if both of the following conditions hold:

- (i) For every term  $t$  in  $D_{c(S)}$ , if “ $t$  is  $F$ ” is true in  $c(S)$  then “ $t$  is  $FH$ ” is true in both  $c$  and  $c(S)$ .
- (ii) For every term  $t$  in  $D_c$ , if “ $t$  is  $FH$ ” is true in  $c$ , then “ $t$  is  $FH$ ” is true in  $c(S)$ .

What condition (i) says, loosely speaking, is that everything in the domain of  $c(S)$  that is  $F$  in  $c(S)$  is  $FH$  in both  $c$  and  $c(S)$ . Satisfaction of condition (i) implies that the set of  $F$ -things in  $D_{c(S)}$  is a subset of the set of  $F$ -things in  $D_c$ . What condition (ii) says, loosely speaking, is that everything in the domain of  $c$  that is  $FH$  in  $c$  is  $FH$  in  $c(S)$ . (This principle of elevation will fail in the case where  $H$  itself contains terms, such as “there”, that may be the products of elevation. But I will ignore that complication.)

In light of this detail regarding the elevation relation, we can understand the truth conditions of sentences (7) and (8). Let us focus just on (8), “Mary said that everyone

working on a graphic interface was present”. I will adjust tenses as needed without further comment. Sentence (8) is true in the context that pertains to the co-worker’s utterance of that sentence, call it  $c$ , because the sentence “Everyone is present” is in the utterance domain that  $c$  assigns to Mary and, as I will explain presently, “Everyone working on a graphic interface was present” is an elevation of “Everyone is present” to  $c$  from  $c$ (“Mary”).

Condition (i) is satisfied, because (we may suppose) the domain of discourse for  $c$ (“Mary”) consists of (terms denoting) members of the committee that she chairs, and for each of them it is true in that context, as well as in  $c$ , that he or she is working on a graphic interface. (I am instantiating “Every  $F$ ” with “every *one*”.) To see the significance of this, suppose that (i) is not satisfied, because there is a member of the domain of the context  $c$ (“Mary”) who is not working on a graphic interface. In that case, (8) would not be true in the context pertinent to the utterance of (8) because (8) would in that case interpret Mary as having said something logically weaker than what she actually said. It would interpret her as having said only that everyone working on a graphic interface was present, when what she actually said was that everyone in a certain more extensive group than that was present. (In some cases, no doubt, but not in this one, we may report a person has having said something logically weaker than what the speaker has said.)

Condition (ii) is satisfied, because (we may suppose) for each person in the domain of  $c$ , which is the context pertinent to the utterance of (8), if it is true in  $c$  that that person is working on a graphic interface, then he or she is also in the domain of  $c$ (“Mary”) and it is true in  $c$ (“Mary”) that he or she is working on a graphic interface. To

see the significance of condition (ii), suppose that it is not satisfied. Suppose that the domain for  $c$  is not included in the domain of  $c$ (“Mary”), so that there are some terms in the domain for  $c$  that are not in the domain of  $c$ (“Mary”). Let  $t$  be one of those terms. Suppose also that it is true in  $c$  that  $t$  is working on a graphic interface. But precisely because  $t$  is not in the domain of  $c$ (“Mary”), it is not true in  $c$ (“Mary”) that  $t$  is working on a graphic interface. (The sentence “ $t$  is working on a graphic interface” will be neither true nor false in  $c$ (“Mary”).) Under these conditions we should not consider (8) to be true either, because it seems to interpret Mary as having said something about a person who is not even (denoted by any term in) the domain of discourse for the context that  $c$  assigns to her; it seems to interpret Mary as having said something about a person working on a graphic interface who was not in fact among the people that Mary was talking about.

I am not taking up here the question of how to define the elevation relation for the case of sentences in which “says that” is embedded under “says that”. If the original utterance by  $Z$  was of the form “Every  $F$  is  $G$ ”,  $X$  may say, “ $Y$  says that  $Z$  says that every  $FH$  is  $G$ ”, but in this case, what amplification  $H$  may be added to the predicate  $F$  will depend on *three* contexts, the context pertinent to  $X$ , the context pertinent to  $Y$ , and the context pertinent to  $Z$ .

In short, my theory about indirect discourse attributions of A-form categorical sentences is that speakers have basically two options. They may simply quote, without modifying the quantifier, but making needed adjustments for words like “here” and “you”. Or they may add modifiers to the quantifier. But in that case the modifier must meet certain conditions with respect to both the context pertinent to the utterance of the indirect discourse sentence and the context that that context assigns to the speaker of the

original utterance. The case in which the indirect discourse sentence precisely subclassifies the universal domain of the original utterance ( $FH = FH \cap D_c = F \cap D_{c(S)}$ ) is a special case of this second alternative.

#### 4. Two kinds of expression

The examples we have considered demonstrate that a distinction may be drawn between two kinds of expression according to how they behave in indirect discourse. On the one hand, there are expressions like “here” and “there”, which I will say *presume wide scope*. These are expressions that need to be rewritten or replaced whenever the context pertinent to the original utterance interprets them differently than the context pertinent to the indirect discourse utterance does. On the other hand, there are expressions like “every” and “everyone”, which I will say *allow narrow scope*. These are expressions that can be left as they were in the original utterance used in the act of speech that the indirect discourse utterance reports on, even when the value of the pertinent context parameter differs between the context pertinent to the original utterance and the context pertinent to the indirect discourse sentence that reports on that utterance.

The distinction between expressions that presume wide scope and expressions that allow narrow scope can be defined in terms of the relation of elevation. To say that an expression  $e$  presumes wide scope is to say that if  $p$  is an elevation of  $q$  to context  $c$  from context  $c(S)$  and  $e$  occurs in  $q$ , then in place of the occurrence of  $e$  in  $q$  we must have in  $p$  an expression  $e'$  that is interpreted relative to the same value in  $c$  that  $e$  is interpreted relative to in  $c(S)$ . For example, “here” presumes wide scope because an elevation of “There is a clown here” must have in place of “here” an expression, possibly “there”, to

which  $c$  assigns the same referent as the context that  $c(S)$  assigns to “here”. On the other hand, to say that an expression  $e$  allows narrow scope is to say that for any pair of sentences  $p$  and  $q$ , though  $p$  is an elevation of  $q$  to  $c$  from  $c(S)$  and  $e$  occurs in  $q$ ,  $e$  may occur *without modification in  $p$* , regardless of the differences between  $c(S)$  and  $c$ .

Expressions that presume wide scope would include personal pronouns (“you” and “I”), temporal indexicals and demonstratives (“now” and “then”), and indexical adjectives (“local” and “recent”). In any conversation, whether or not such expressions occur in the “that”-clause of an indirect discourse sentence uttered in that conversation, their interpretation is that which is assigned to them *directly* by the context that pertains to the conversation, not the interpretation that might be assigned to them *indirectly* by the context that that context assigns to the speaker who made the original utterance used in the speech act that the utterance of the indirect discourse sentence reports on.

(Alterations in tense may be due not to the difference between the context pertinent to the original utterance and the context pertinent to the indirect discourse utterance but due to the tense of “say” and sequence-of-tense rules.)

Expressions that allow narrow scope include expressions that are explicitly or implicitly quantificational. Explicitly quantificational expressions are quantifiers like “every”, “some” and “most”. Implicitly quantificational expressions are those such that when we formulate the truth conditions of sentences formed from them, we do so by quantifying over appropriate entities of some kind. “Might” and “must” are implicitly quantificational in this sense because we explicate “might” with an existential quantification over possibilities of some kind and explicate “must” with a universal

quantification over possibilities of that kind. Another example might be conditional operators such as “if . . . then . . .”

I do not have a general explanation of *why* some expressions presume wide scope and others allow narrow scope. The fact that the expressions that presume wide scope are in some way demonstrative while the expressions that allow narrow scope are in some way quantificational provides perhaps a clue. In the case of demonstratives that are also indexicals, there are some porous semantic rules, such as that “here” refers to the place at which the utterance takes place. (The rule does not fix an interpretation, since it says nothing about the size of the region, and it is “porous” because there are exceptions, such as when one puts one’s finger on a map and says “here”.) In the case of demonstratives that are not indexicals, such as “this” and “that”, we can interpret distinct occurrences of those demonstratives differently within a single context (as when we say “*This* is bigger than *this*”) by exploiting a variety of external cues, such as pointing and parallel structures elsewhere in “the text” (Gauker 2008). The states of affairs that the rules for indexicals direct us to always pertain immediately to the utterance being interpreted, not some remote time or place. And the cues that we go by in deciding how to interpret a nonindexical demonstrative are always cues surrounding the execution of the current utterance.

By contrast, our means of identifying the contextually determined domains of quantification pertinent to the interpretation of in-some-way quantificational expressions is not as tightly constrained by conventions. Here we usually have little more to go by than the consideration of what domain is most relevant under the circumstances of utterance (Gauker 1997). Since we are often in no position to identify the set of things

that are relevant in this way to the act of speech being reported on, we allow ourselves the option of, in effect, simply quoting these expressions and leaving it just unclear what domain was pertinent to the act of speech being reported on. And if we are not going to simply quote but are going to try to capture, from the point of view of the context of attribution, what the speaker was saying in the different context in which he or she spoke, then because we have no conventional means of indicating any particular domain, we have to choose our words in a way that takes into account both the domain of the context pertinent to the attribution and the domain of the context pertinent to the speaker's original utterance.

### **5. The paradox of indirect discourse**

A semantics for indirect discourse ought to be able to tell us what is wrong with the following argument:

*The paradox of indirect discourse*

1. Silas says that everybody is wearing a funny hat. (Premise)
2. What Silas says is true. (Premise)
3. What Silas says is that everybody is wearing a funny hat. (From 1)
4. That everybody is wearing a funny hat is true. (From 2 and 3, by substitution of identicals)
5. Everybody is wearing a funny hat. (From 4, by semantic descent)

This a "paradox" because, while the premises "appear to be true in the context in which Emma is speaking, the conclusion is false in that context, and yet each of the inferences appears to be valid in the sense that it preserves truth-in-a-context. My diagnosis of the



error in this argument will not turn on the details of my semantics for indirect discourse, but I want to say something about this because relativists might offer a diagnosis of this argument in support of their own analysis.

The relativist diagnosis will be that Premise 2 is false. This is what EHW and Egan say, in effect (EHW, p. 145; Egan 2007, p. 3), although they do not consider the paradox in exactly this form. (EHW pose a very similar paradox, pp. 133-134). What the subject term of Premise 2 denotes, they in effect say, is Silas's *utterance*, which from the point of view of Emma's context of evaluation is false. My answer to this is that they are wrong in asserting that from Emma's point of view Silas's utterance is false. Silas's utterance is *true simpliciter*, since the sentence he utters is true in the context pertinent to his utterance, and so *a fortiori* it is true from every point of view. What we can say is that the *sentence* Silas utters would not be true in the context pertinent to Emma's conversation; so if what Premise 2 means is that the *sentence* Silas uttered is true in Emma's context, then Premise 2 is false.

Egan acknowledges (2005, p. 3, note 3) that someone might diagnose his evaluation (of his analogue to (2)) as the product of mistaking the falsehood of Silas's *sentence* in Emma's context for the falsehood of Silas's *utterance* of that sentence. His answer is that people can perfectly well distinguish between evaluating the truth value of an utterance and evaluating the truth value of the sentence uttered relative to their own context. But having said that, he does nothing to show that our ability to draw that distinction would be misapplied in evaluating "what Silas said" as referring to Silas's utterance and evaluating that utterance as true.

There might be some temptation to say that the conclusion of the argument is not actually false relative to Emma's context. Not everybody in the domain of discourse pertinent to Emma's context is wearing a funny hat, to be sure. But in view of its provenance, as a conclusion from the argument that precedes, it might be said that even in Emma's situation, the sentence has to be evaluated with respect to "projected" values of the contextual parameters, namely, the parameters determined by Silas's situation (cf., EHW, p. 162; Weatherson 2008, p. 536). However, I think this answer is indefensible. For a case in which parameters may indeed be projected in this way, suppose Emma utters the following sentences:

(9) Silas says that there is a clown there. Everybody is wearing a funny hat.

In (9), we can take Emma's utterance of "Everybody is wearing a funny hat" to be part of what she says Silas says and evaluate her utterance by evaluating the sentence uttered relative to the parameters of the context pertinent to Silas. But when we are evaluating the validity of the argument from 1 and 2 to 5 in the paradox of indirect discourse, we should not evaluate the conclusion relative to a context different from the one relative to which we evaluate the premises, and the evaluation of the conclusion in that context does not call for appealing to the values that the contextual parameters have according to some other context (such as the one pertaining to Silas).

Here is what is really wrong with the argument. If 2 means that the *sentence* Silas uttered is true in Emma's context, then 2 is false, and, moreover, on that reading 3 does not even make sense. If 2 means that Silas's *utterance* is true, then 2 may be true in Emma's context, but in that case the "is" in 3 is not the "is" of identity but the "is" of predication, and what 3 says about that utterance is that it belongs to a certain type, the

type of utterance characterizable as *that everybody is wearing a funny hat*. But in that case, 4 does not follow from 2 and 3 by substitution of identicals or by any other valid rule of inference.

Someone might think that yet another possible meaning of 2 is that the *proposition* that Silas expresses is true. That is the reading on which 3 might look like a true identity statement. However, I myself have not explicated the “that”-clauses of indirect discourse sentences as denoting propositions, and I do not see any persuasive reason to do so. (Various sources of the proposition idea are identified and dismissed in my 2003.) I do not deny that there is a perfectly good use for the term “proposition” in ordinary language. But that fact should not encourage us to read 2 in this way either. If it is ordinary language that warrants this reading, then, by the same token, ordinary language should warrant us in asserting, “What Silas says is the proposition that Silas expresses.” But as a piece of ordinary language, that sentence does even not quite make sense.

## **6. The basis for a semantics for indirect discourse**

A commonplace conception of the task of semantic theory is that it has to account for our linguistic understanding. That is, it has to explain what it is we understand about a language that puts us in a position to understand any given utterance of that sentence. My own view is that this is not a very satisfactory characterization of the job of a semantic theory, because it begs the question: What is it that we have to understand in order to understand a sentence?

A partial conception that does not beg the question is that part of the task of a semantic theory for an expression is to enable us to define the class of logically valid arguments involving that expression. (This was the primary criterion by which I defended a semantics for conditionals in my 2005.) In the case of indirect discourse sentences, however, this test does no work since the “that”-clauses are logically completely inert. If  $p$  and  $q$  are different sentences, then, whatever the logical relation between them, there is simply no logical relation, other than logical consistency, between “ $S$  says that  $p$ ” and “ $S$  says that  $q$ ”. They will not be logically inconsistent, and neither logically implies the other. Perhaps the sentence “Silas says that a clown and a pony are there” logically implies “Silas says that a clown is there”, but even that is questionable. Going by logic alone, then, we could not have any reason to treat the denotation of a “that”-clause in an indirect discourse sentence as anything other than an unstructured atom distinct from the denotation of every other “that”-clause in an indirect discourse sentence but not differentiated from those other atoms in any particular way.

Another conception of the task of semantic theory might be that it has to tell us how sentences correspond to reality when they are true. In the case of indirect discourse sentences, then, we would need to identify the sorts of states of affairs in the world that would make an indirect discourse sentence true and try to explain how the elements of the indirect discourse sentence relate to the elements of those states of affairs. For instance, we might think that *saying* is a three-place relation between speaker, a sentence (or utterance) and a proposition and try to build all three of those into our account of the truth conditions of an indirect discourse sentence. The main problem with this approach is that it presumes that we understand the denotation relation, which supposedly holds

between words and things (relative to a context), whereas, as a matter of fact, despite a long effort, no one has ever been able to give a plausible and reasonably comprehensive account of what that relation might be.

There is, however, at least one other possible kind of datum that we can put to use in defending a semantics for indirect discourse, namely, the kind of datum that I have employed in this essay concerning the way in which the “that”-clause by means of which we report an act of speech has to differ from the sentence uttered in the act of speech on which we are reporting. For example, from the fact that Emma has to substitute “there” for “here” when she reports Silas’s act of speech in (3), we learn that our semantic account has to say something specifically about the interpretation of “there” as it occurs in the “that”-clause. In particular, our semantics has to provide for the possibility that the interpretation of “there”, as it occurs in the “that”-clause, is identical to the interpretation of “here”, as it occurs in Silas’s original utterance. And from the fact that we can substitute “everyone who is working on a graphic interface” in the “that”-clause in (8) for the bare “everyone” in Mary’s original utterance, we learn that our semantic account has to provide for a domain of discourse for the sentence in the “that”-clause that is potentially distinct from the domain of discourse for the indirect discourse sentence itself.

These sorts of data are special cases of a more general kind of datum that we can appeal to in defending a semantic theory. Roughly speaking, these data have to do with the ways in which we can use language to *transmit information*. It is not only in indirect discourse that we may be called upon to alter a speaker’s own words. We also have to do this when we accept another person’s testimony and go to pass it on to someone else. Instead of telling Dottie what Silas said, Emma might have decided to accept Silas’s

assertion and make the same assertion for herself. However, in that case, she could not have said to Dottie, “Everyone is wearing a funny hat.” In that case, she would have had to know enough about the context of Silas’s utterance in order to modify it in way that would make clear to Dottie whom she was talking about. So she might have said, “Everyone at Jonas’s party is wearing a funny hat”. Indeed, indirect discourse might be conceived as a special case of such episodes of information transmission. It is the special case in which the speaker ascribes responsibility for the saying to the previous speaker rather than accept responsibility for it him- or herself. The more general area of concern, which a semantic theory can address, is the ways in which a sentence may or must be transformed when it is grounded in an utterance to which a different context pertains than the context that pertains to the present conversation.

### **7. Against a pragmatic theory**

In part, our question has been: In what ways may the “that”-clause in an indirect discourse sentence elaborate on the words that the attributee actually spoke? Many people, confronted with this question, would be inclined to answer as follows: We may add to the “that”-clause whatever words we need to add in order to *clarify the speaker’s meaning*. But that, I now wish to argue, while fine as folk linguistics, is a flatly question-begging answer, if taken as theoretically fundamental.

First of all, not everything a speaker *means* by what he or she says is something he or she *says*. So though in saying “It’s getting a bit chilly in here, don’t you think?” you may *mean* that I should get up and close the window, we should not characterize you as having *said* that I should close the window. Rather, we should say that you *implied* as

much. But I am not sure that we could draw the needed distinctions without switching to some kind of truth-conditional account such as I have offered here in terms of contexts that point to other contexts. In getting clear about the kind of meaning expressed in the “that”-clauses of indirect discourse sentences, we may find ourselves undertaking precisely the kind of semantic project that I have been executing in this paper.

What else, if not the kind of project I am engaged in here, could tell us the ways in which in indirect discourse we may “clarify” what the speaker *said*? The thought might be that the speaker had a definite proposition in mind and that what it takes to express that proposition in the situation that the attributor is in is different from what was required of the speaker in order to express that proposition in the situation in which the speaker spoke. So what we have to do, in choosing the wording of the “that”-clause in an indirect discourse sentence, is find words that express that same thought in the new situation that the speaker’s own words expressed in his or her own situation.

Thus we are led to the question, how does a speaker manage to express a thought using words, and how do we interpreters manage to express that thought in our own situation? Much contemporary theorizing concerning the nature of linguistic communication unrealistically presupposes that interlocutors have the ability to detect one another’s intentions and other states of minds independently of the interpretation of what they say so that they may then appeal to those states of mind in identifying the thoughts that the speaker’s words express. I have criticized this presupposition in a number of prior publications (1997, 2001, 2003, 2008), and I will not rehearse those criticisms here. The moral I draw is that our access to people’s thoughts is primarily an understanding of what they *say* that does not depend on a prior understanding of what

they have in mind. Normally, the hearer can do no more than try to understand what the speaker has said independently from any assumptions about what the speaker has in mind. Normally, the question of what the speaker has in mind, potentially distinct from what he or she says in this sense, does not even arise. When it does arise, if the answer is not just that what the speaker has in mind is the same as what he or she said, then the process of discerning that difference will begin with an independent identification of what the speaker has said.

On such grounds I conclude that clarifying what the speaker said can only mean finding a sentence such that the conditions under which it is true-in-a-context are suitably related to the conditions under which the sentence the speaker uttered is true-in-a-context. The partial account of the elevation relation in section 3 above illustrates an approach to explicating precisely this relation. Thus we are led back to the kind of semantic theory that I have been developing in this paper.

These observations bring to light an important fact about the concept of context that I have been employing here: Contexts have to be thought of as something objective. That is, what belongs to the context pertinent to an utterance is not determined by what the speaker thinks it is. Otherwise the hearer would be in the position of having to read the speaker's mind in order to understand what the speaker has said, which I have claimed is not usually possible. Whereas, if the content of the context is determined objectively by the state of the world in which the conversation takes place, the hearer can (fallibly) use his or her knowledge of that state of the world to draw conclusions about the pertinent context and on that basis draw conclusions about what the speaker has said.



If, as I suppose, but have not argued here, the truth of an utterance depends on *nothing but* the content of the context that pertains to that utterance (as well, of course, on semantic facts concerning the conditions a context must meet in order for the sentence uttered to be true in it), then there is another reason too to suppose that the content of the pertinent context is an objective matter. We think that speaking the truth is something a speaker has to strive to do and can succeed in doing only by taking account of the way the world is around him or her. But if truth depended on nothing but context and context depended entirely on what the speaker had in mind, then speaking the truth would not be something that required that kind of striving. (The toy semantics in the appendix will embody my assumption that the truth of a sentence in a context depends on nothing but the content of the context.)

Though the context that pertains to a conversation is objective, and not determined by the states of mind of speaker or hearer, a speaker's choice of words may reflect the speaker's representation of the hearer's representation of the context, when there is some basis for making assumptions about this. We have seen that indirect discourse sentences may approximate to quotations. Thus Emma may say to Dottie, (5) "Silas says that everyone is wearing a funny hat" – that is, that sentence may be true in the context pertaining to Emma's conversation with Dottie — even if Dottie does not know where Silas is. But if Emma knows that Dottie does not know where Silas is — that is, Emma's representation of Dottie's representation of the context pertinent to their conversation shows that Dottie does not represent that context as assigning to Silas a context in which the domain of discourse is people at Jonas's party — then she will have reason to say, instead, (6) "Silas says that everyone *at Jonas's party* is wearing a funny

hat”. The latter will do more than the former to enable Dottie to correctly represent the context that pertains to their conversation (inasmuch as that context assigns to Silas a context pertinent to his utterance).

None of this is an objection to the claim that in formulating an indirect discourse sentence we need to choose the words of the “that”-clause in a way that clarifies what the speaker said. That is a perfectly fine piece of folk linguistics. We might help our children formulate their indirect discourse sentences by telling them to choose their words in a way that clarifies what the attributee said. But that does not mean that we theorists can avoid the sort of semantic theory of indirect discourse that I have been engaging in here. This is just the kind of theory we will need to develop when we go to explain in detail what “clarifying what the said” really amounts to.

### **8. Extending the account to other propositional attitude sentences**

The present semantics for indirect discourse does not in any obvious way extend to other sorts of sentences containing “that”-clauses, such as sentences that attribute beliefs. The problem in the case of belief is that when we attribute a belief our attribution may not be grounded in any overt utterance in the way in which an indirect discourse statement may be.

However, in view of the *way* in which the account of indirect discourse depended on the attributee’s having uttered something, I think the obstacles may not be so high. A context, as here defined, assigns to each of a number of people (strictly speaking, to their names) an *utterance* domain. The utterance domain for a person is a set of *sentences*. It is not a set of *utterances*. Actual utterances come into the account only insofar as we

may expect that in order for a context to be the one that pertains to a given conversation, the sentences in the utterance domain for a speaker must be sentences that that speaker has actually uttered. So similarly, in formulating a semantics for belief sentences, we could suppose that a context assigns to each of the people to whom beliefs can be attributed a set of sentences, the *belief set* for that person, representing that person's *beliefs*. In this case, the requirement that a context would have to meet in order for the context to pertain to a given conversation would not be that each of the sentences in the belief set for a person be a sentence that that person actually have uttered, but only that for a certain other relation, the person stand in that relation to each of the sentences in the set.

The big question will be: What is the relation that a person must stand in to each of the sentences in the belief set that a context assigns to that person in order for the context to pertain to a conversation? One possibility would be that each of those sentences is a "direct translation" of a sentence written in the "belief box" in the person's brain. Within this option we could distinguish different theories about the nature of the language of brain writing. (It could be a form of the same language that the believer speaks, or it could be a special kind of mentalese.) More precisely, to allow for inexplicit beliefs, we might say that it is a translation of a logical consequence of sentences written in the belief box. A different possibility would be to say that the sentences in the belief set for a person are sentences that that person would be disposed to speak if he or she were asked to state his or her beliefs and had no fear of retribution or other reason to dissemble. Probably neither of these answers is correct quite as it stands. But they

perhaps give us hope of being able to formulate a semantics for belief sentences analogous to the present semantics for indirect discourse sentences.

I should emphasize that such an account of the semantics of belief attributions would *not* entail analyzing *belief* as a relation between a person and a sentence. The present approach to the semantics of “believes” does not proceed by analyzing the relation of belief and then assigning that relation to the word “belief” as its denotation.

### 9. The *de dicto/de re* distinction

So far I have not had anything to say about the *de dicto/de re* distinction, which is usually a large part of what is at issue in philosophical discussions of indirect discourse and attributions of propositional attitudes. Here I will merely state a couple of opinions on the subject without attempting to defend them.

Most philosophers and semanticists (but not all) will grant that the components of “that”-clauses in indirect discourse sentences are normally referentially opaque in the sense that substitution of co-extensional phrases does not preserve truth. However, it is often supposed that there is a special kind of indirect discourse sentence, the *de re* kind, in which one or more of the components is referentially transparent. So we may regard “*S* said that *a* is *F*” and “*a* = *b*” as logically implying “*S* said that *b* is *F*”, and in that case, “*S* said that *a* is *F*” is said to be *de re*, as opposed to *de dicto*. Moreover, it is often supposed that there is a special syntactic form that can be used to make it explicit that an indirect discourse statement is *de re*. If we say, “*S* said of *a* that it [he/she] is *F*”, then that is supposed to show that the occurrence of “*a*” is referentially transparent. It does so by placing the referentially transparent term outside of the “that”-clause.

The first opinion I would like to record is that the *de re* indirect discourse statement is an entirely fictitious beast. “*S* said that *a* is *F*” and “*a* = *b*” never logically imply “*S* said that *b* is *F*”, not even when the supposedly dedicated syntax is used. That is, “*S* said of *a* that he is *F*” and “*a* = *b*” does not imply “*S* said of *b* that he is *F*”. The “said of *a* that” locution is merely a device by which we may put the focus on *a* (as I said in my 2003, pp. 269-70). Consider the following dialogue:

*Mom*: Billy said that Santa Claus would bring him a bicycle.

*Dad*: Billy doesn't think we're giving him a bicycle.

*Mom*: No, Billy said of *Santa Claus* that he would bring him a bicycle.

Mom, in her second statement, clearly is not committing herself to the existence of Santa Claus. She is merely putting the focus on “Santa Claus” in order to emphasize to Dad that it is not they, the parents, who Billy thinks will bring him a bicycle.

Though “*S* said that *a* is *F*” and “*a* = *b*” do not logically imply “*S* said that *b* is *F*”, there are cases in which we are prepared to *accept* both “*S* said that *a* is *F*” and “*S* said that *b* is *F*” and in which the truth of “*a* = *b*” is part of the reason why both are acceptable. And that fact may be part of the reason why some people believe in the existence of *de re* indirect discourse statements. The second opinion that I would like to record is that these cases may be treated as cases in which “*b* is *F*” is an elevation, in my sense, of “*a* is *F*” to the context pertinent to the utterance of “*S* said that *b* is *F*” from the context pertinent to *S*'s utterance of “*a* is *F*”. What I would need to do next is work out the conditions that a context must meet in order for this elevation relation to hold in it, but I will not try to do that here.

## Appendix

The purpose of this appendix is to describe precisely a semantics for a simple language permitting indirect discourse. I will use a sans-serif font both for expressions in the object language and for metalinguistic variables ranging over such expressions. I will “leave it to context”, as they say, to distinguish which is which.

### *Syntax*

First, we define a language  $L$  without **says that** and, then, in terms of that, define a language  $L^+$  containing **says that**. The syntax of the language  $L$  will be like the usual syntax of the languages of logical studies. Atomic formulae will be formed from predicates and individual variables and singular terms in the usual way. The singular terms of  $L$  include **here** and **there**. The two-place predicates of  $L$  include the identity sign  $=$ . Compound formulae will be built up from atomic formulae, negation symbols, disjunction symbols and parentheses in the usual way. However, I will assume that all quantifications have the form:  $\forall x(F : G)$ , where  $F$  and  $G$  are formulae of  $L$ . *Sentences* are formulae with no free variables. Every sentence of  $L$  is a sentence of  $L^+$ , and if  $n$  is a singular term of  $L$  and  $p$  is a sentence of  $L$ , then  $n$  **says that**  $p$  is a sentence of  $L^+$ .

(Since  $p$  has to be a sentence of  $L$ , **says that** will not be embedded under **says that**.) (**F & H**) abbreviates  $\neg(\neg F \vee \neg G)$ . If  $p$  is a sentence and  $n$  and  $m$  are singular terms or individual variables of  $L$ , then  $p[n/m]$  is the result of substituting an occurrence of  $n$  for every occurrence of  $m$  in  $p$ .

## Contexts

The set of contexts for  $L+$  will be defined recursively. For the basis, we define *basic* contexts as contexts that assign only the empty set as the context determined for an agent by the context. (These are contexts relative to which, as it were, nobody says anything.) Then in terms of these we define contexts proper. Say that two singular terms  $c$  and  $d$  are *identity-linked* in a set  $S$  if and only if either  $c = d \in S$  or there is a term  $e$  such that  $c$  is identity-linked to  $e$  in  $S$  and  $e$  is identity-linked to  $d$  in  $S$ . Where  $S$  is a set the members of which are either atomic sentences of  $L$  or negations of atomic sentences of  $L$ , say that  $S$  is a *consistent set of literals* if and only if (a) there is no sentence  $p$  such that both  $p$  and  $\neg p$  are in  $S$ , and (b), where  $v_1, v_2, \dots, v_n$  are individual variables in  $p$ , if for each  $i$ ,  $1 \leq i \leq n$ ,  $c_i$  and  $d_i$  are identity-linked, then not both  $p[c_1/v_1] \dots [c_n/v_n]$  and  $\neg p[d_1/v_1] \dots [d_n/v_n]$  are in  $B_\Gamma$ .

A *basic context*  $\Gamma = \langle B_\Gamma, N_\Gamma, S_\Gamma, \delta_\Gamma, \sigma_\Gamma \rangle$ , where

$B_\Gamma =$  (the base of  $\Gamma$ ) a consistent set of literals (not necessarily maximal),

$N_\Gamma =$  (the domain for  $\Gamma$ ) a set of singular terms containing at least every singular term that occurs in any member of  $B_\Gamma$ ,

$S_\Gamma =$  (the speaker domain for  $\Gamma$ ) a set of names (names of people to whom discourse may be attributed in  $\Gamma$ ),

$\delta_\Gamma =$  a function that assigns to each member of  $S_\Gamma$  a set of sentences of  $L$ , and

$\sigma_\Gamma =$  a function that assigns the empty set to each member of  $S_\Gamma$ .

A *context*  $\Gamma = \langle B_\Gamma, N_\Gamma, S_\Gamma, \delta_\Gamma, \sigma_\Gamma \rangle$ , where  $B_\Gamma$ ,  $N_\Gamma$ ,  $S_\Gamma$ , and  $\delta_\Gamma$  are defined as before,

but

$\sigma_\Gamma =$  a function that assigns to each member of  $S_\Gamma$  either a basic context or a context.

(See figure 1.) (Since **says that** is never embedded under **says that** in  $L^+$ , we could confine our attention to contexts that assign only *basic* contexts to the names in the speaker domain.)

In order to avoid writing subscripts on subscripts, I will sometimes omit subscripts. Thus it may be understood that  $B_{\sigma(n)}$  is the base of the context that  $\sigma_{\Gamma}$  assigns to  $n$ .




Figure 1: A schematic picture of a context

### *The elevation relation*

The elevation relation is a four-place relation that holds between two sentences and two contexts. We define it in two stages. First, we define several different kinds of elevation,



one for each modification that may have to be made. Then we define the elevation relation as the product of a sequence of partial elevations.

$p$  is a *here-only-elevation* of  $q$  to  $\Gamma$  from  $\sigma_\Gamma(n)$  if and only if either:

- (a)  $q$  contains **here** but not **there** and  $p = q$ , and for some singular term  $t$  of  $L$ ,  $t = \text{here}$  (or  $\text{here} = t$ ) is a member of both  $B_{\sigma(n)}$  and  $B_\Gamma$ , or
- (b)  $q$  contains **here** but not **there** and  $p$  is the result of substituting **there** for every occurrence of **here** in  $p$ , and for some singular term  $t$  of  $L$ ,  $t = \text{here}$  (or  $\text{here} = t$ ) is a member of  $B_{\sigma(n)}$  and  $t = \text{there}$  (or  $\text{there} = t$ ) is a member of  $B_\Gamma$ .

$p$  is a *there-only-elevation* of  $q$  to  $\Gamma$  from  $\sigma_\Gamma(n)$  if and only if either . . . (similarly, but with **here** and **there** reversed).

$p$  is a *here-and-there-elevation* of  $q$  to  $\Gamma$  from  $\sigma_\Gamma(n)$  if and only if either:

- (a)  $q$  contains both **here** and **there**, and  $p = q$ , and for some singular term  $t_1$  of  $L$ ,  $t_1 = \text{here}$  (or  $\text{here} = t_1$ ) is a member of both  $B_{\sigma(n)}$  and  $B_\Gamma$ , and for some singular term  $t_2$  of  $L$ ,  $t_2 = \text{there}$  (or  $\text{there} = t_2$ ) is a member of both  $B_{\sigma(n)}$  and  $B_\Gamma$ , or
- (b)  $q$  contains **here** and **there** and  $p$  is the result of simultaneously substituting **there** for every occurrence of **here** and **here** for every occurrence of **there** in  $q$ , and for some singular term  $t_1$  of  $L$ ,  $t_1 = \text{here}$  (or  $\text{here} = t_1$ ) is a member of  $B_{\sigma(n)}$  and  $t_1 = \text{there}$  (or  $\text{there} = t_1$ ) is a member of  $B_\Gamma$ , and for some singular term  $t_2$  of  $L$ ,  $t_2 = \text{there}$  (or  $\text{there} = t_2$ ) is a member of  $B_{\sigma(n)}$  and  $t_2 = \text{here}$  (or  $\text{here} = t_2$ ) is a member of  $B_\Gamma$ .

For each expression of the form  $\forall x(F: \text{ of } L, \forall x((F \ \& \ H):$  is an *expansion* of  $\forall x(F:$  relative to  $\Gamma$  and  $\sigma_{\Gamma}(n)$  if and only if:

- (a) for every term  $t$  in  $N_{\sigma_{\Gamma}(n)}$ , if  $F[t/x]$  is true in  $\sigma_{\Gamma}(n)$ , then  $(F \ \& \ H)[t/x]$  is true in both  $\Gamma$  and  $\sigma_{\Gamma}(n)$ , and
- (b) for every term  $t$  in  $N_{\Gamma}$ , if  $(F \ \& \ H)[t/x]$  is true in  $\Gamma$ , then  $(F \ \& \ H)[t/x]$  is true in  $\sigma_{\Gamma}(n)$ .

$p$  is a *quantifier-elevation* of  $q$  to  $\Gamma$  from  $\sigma_{\Gamma}(n)$  if and only if:

- (a)  $q$  contains an expression of the form  $\forall x(F:$ , and
- (b)  $p$  is the result of replacing each expression of the form  $\forall x(F:$  in  $q$  with an expansion of that expression relative to  $\Gamma$  and  $\sigma_{\Gamma}(n)$ .

The here-only-, there-only-, and here-and-there-elevations are *obligatory* elevations, and the quantifier-elevation is an *optional* elevation.

Now we can define the elevation relation for  $L+$  as follows:  $p$  is an *elevation of  $q$*  to  $\Gamma$  from  $\sigma_{\Gamma}(n)$  if and only if there is a sentence  $s$  such that (a) either  $q$  contains no occurrence of **here** or **there** and  $q = s$ , or  $s$  is an obligatory elevation of  $q$  to  $\Gamma$  from  $\sigma_{\Gamma}(n)$ , and (b) either  $s = p$  or  $p$  is an optional elevation of  $s$ . In other words, we obtain an elevation of  $q$  to  $\Gamma$  from  $\sigma_{\Gamma}(n)$  by applying to  $q$  whatever obligatory elevation applies and then optionally applying the quantifier-elevation. The definitions are written in such a way that if a sentence contains more than one expression of the form  $\forall x(F:$ , then if any one is expanded, then all of them must be expanded.

For example,  $\forall x((F \ \& \ H): R(x, \text{there}))$  may qualify as an elevation of  $\forall x(F: R(x, \text{here}))$  to  $\Gamma$  from  $\sigma_\Gamma(n)$ , because  $\forall x(F: R(x, \text{there}))$  may be a here-only-elevation of  $\forall x(F: R(x, \text{here}))$  to  $\Gamma$  from  $\sigma_\Gamma(n)$ , and  $\forall x((F \ \& \ H): R(x, \text{there}))$  may be a quantifier-elevation of a sentence of the form  $\forall x(F: R(x, \text{there}))$  to  $\Gamma$  from  $\sigma_\Gamma(n)$ .

### *Truth conditions*

Some sentences will be neither true nor false in some contexts. So we will provide separate formulations of truth and falsehood conditions for sentences of  $L^+$ .

- (T0) If  $p \in B_\Gamma$ , then  $p$  is true in  $\Gamma$ .
- (T $\neg$ ) If  $p$  is false in  $\Gamma$ , then  $\neg p$  is true in  $\Gamma$ .
- (T $\vee$ ) If  $p$  is true in  $\Gamma$  or  $q$  is true in  $\Gamma$ , then  $(p \vee q)$  is true in  $\Gamma$ .
- (T $\forall$ ) If, for every  $n \in N_\Gamma$  such that  $F[n/x]$  is true in  $\Gamma$ ,  $G[n/x]$  is true in  $\Gamma$ , then  $\forall x(F: G)$  is true in  $\Gamma$ .
- (TID) If for some sentence  $q \in \delta_\Gamma(n)$ ,  $p$  is an elevation of  $q$  to  $\Gamma$  from  $\sigma_\Gamma(n)$ , then  $n$  says that  $p$  is true in  $\Gamma$ .
- (TCI) No other sentence is true in  $\Gamma$ .
- (F0) If  $\neg p \in B_\Gamma$ , then  $p$  is false in  $\Gamma$ .
- (F $\neg$ ) If  $p$  is true in  $\Gamma$ , then  $\neg p$  is false in  $\Gamma$ .
- (F $\vee$ ) If  $p$  is false in  $\Gamma$  and  $q$  is false in  $\Gamma$ , then  $(p \vee q)$  is false in  $\Gamma$ .
- (F $\forall$ ) If for some  $n$  such that  $F[n/x]$  is true in  $\Gamma$ ,  $G[n/x]$  is false in  $\Gamma$ , then  $\forall x(F: G)$  is false in  $\Gamma$ .

(FID) If for every sentence  $q \in \delta_r(n)$ ,  $p$  is not an elevation of  $q$  to  $\Gamma$  from  $\sigma_r(n)$ , then  
n said that  $p$  is false in  $\Gamma$ .

(FCI) No other sentence is false in  $\Gamma$ .

Note that in (F $\forall$ ) we might have written “for some  $n \in N_r$ ”, but in view of the stipulation that  $N_r$  contain every singular term in any member of  $B_r$ , doing so would have been redundant.

### Acknowledgements

This research was supported by grants from the Taft Research Center of the University of Cincinnati. A discussion with Michael Glanzberg in Oslo provoked me to get clearer about some of the matters discussed in this paper.

### References

- Egan, Andy, John Hawthorne and Brian Weatherson, 2005: “Epistemic Modals in Context,” in G. Preyer and G. Peter, eds., *Contextualism in Philosophy: Knowledge, Meaning and Truth*, Oxford University Press, pp. 131-168.
- Egan, Andy, 2007: “Epistemic Modals, Relativism and Assertion,” *Philosophical Studies* 133: 1-22.
- Gauker, Christopher, 1997: “Domain of Discourse,” *Mind* 106: 1-32.
- Gauker, Christopher, 2001: “Situated Inference versus Conversational Implicature,” *Noûs* 35: 163-189.
- Gauker, Christopher, 2003: *Words without Meaning*, MIT Press.

Gauker, Christopher, 2005: *Conditionals in Context*, MIT Press.

Gauker, Christopher, 2008: "Zero Tolerance for Pragmatics," *Synthese* 165: 359–371.

Weatherson, Brian, 2008: "Attitudes and Relativism," in John Hawthorne, ed.,  
*Philosophical Perspectives*, Vol. 22, *Philosophy of Language*, John Wiley & Sons:  
527-544.