

Bilboards, Bombs and Shotgun Weddings¹

Andy Egan

Australian National University/University of Michigan

Draft of June 22, 2007

Introduction

A lot of natural language sentences are *context sensitive* – the way that they represent things as being, and the message that their use conveys to their audience, depends on the context in which they're used. When Butch says "I am hungry", he expresses the proposition *that Butch is hungry*. When Sundance says, "I am hungry", he expresses the proposition *that Sundance is hungry*. Somewhat more controversially: in some contexts, the sentence "Sundance is ready" expresses the proposition *that Sundance is ready for breakfast*. In other contexts, the sentence "Sundance is ready" expresses the proposition *that Sundance is ready to lay down covering fire*, or *that Sundance is ready to make a break for it*.

There is a standard and familiar story about how to model this sort of context dependence in our semantic theorizing: Sentences aren't associated with contents *simpliciter*, but with characters, which are functions from contexts of utterance to contents. A sentence's character captures the way in which how it represents, what message an utterance of it conveys to the members of its audience, depends on the context in which it is used. A sentence has the character it has because its component expressions have the characters *they* have. Lexical items don't (in general) have semantic values *simpliciter* – they have characters, which are functions from contexts of utterance to semantic values. A lexical

¹ Thanks to Hud Hudson, Theodore Sider, Frank Jackson, David Chalmers, Avery Andrews, Josh Parsons, Joshua Brown, Timothy Sundell, David Braddon-Mitchell, Daniel Nolan, Thony Gillies, James Joyce, Nicholas Silins, Kent Bach, and David Velleman, and to audiences at the Australasian Association of Philosophy annual conference, the Australian National University, and the University of Canterbury, the University of Wyoming, as well as to the RSSS contextualism reading group, for extremely helpful questions, comments, and objections. I'm also indebted to a series of papers by Stefano Predelli (1996, 1998a, 1998b), and a paper by Alan Sidelle (1991), which I was directed to after this paper was mostly complete. The story I'll run here has some obvious affinities with the one proposed in Predelli 1996 – I take it to be in competition with Predelli's view as an explanation of an overlapping class of phenomena. I regret that restrictions of space and time prevent me from engaging more directly with his proposal here.

item's character captures the way that the contribution that it makes to the content of the sentences in which it occurs depends on the context in which it is used.

What's a *context*, exactly? One very natural thing to say is that "a context is a location – time, place, and possible world – where a sentence is said..." (Lewis 1980, p.21). On this sort of picture, a context is something like a potential point of origin for an utterance – it's a position or situation in which an utterance does or could (in some suitably extended sense of 'could') occur.² What we want a context to do is fix the semantically relevant properties of the speaker, or of the speaker's particular situation, position, or predicament.³ Classes of such possible predicaments – or anyway, classes of things that are well-suited to model such predicaments, such as centered worlds – look like ideal candidates for doing this. (It's easy to read Kaplan 1989 as endorsing this sort of picture as well.) It's very natural to follow Lewis (1980, p.28) in thinking that, once we've fixed where (and therefore who) the speaker is, which world the utterance takes place in, and the time at which it occurs, we'll have fixed all of the facts about the utterance on which its content could depend. Fix that the world of utterance is *w*, the time of utterance is *t*, and the speaker of the utterance is *x*, and you'll have fixed all of the speaker's communicative intentions: they're the communicative intentions that *x* has at *t* in *w*. You'll have fixed what's salient: it's whatever's salient to *x* (or to those with whom *x* is engaged in conversation) at *t* in *w*. You'll have fixed the standards of precision: they're the standards that are in effect in the conversation to which *x* is a party at *t* in *w*. And so on.

² We don't want to restrict ourselves to contexts in which the sentence is actually uttered, since for various purposes it will be useful to assign contents to sentences relative to contexts in which they're not uttered. This is not a problem - we can still specify the features that content's sensitive to, whether the utterance occurs or not.

³ Perhaps the most natural word to use for the things I want to talk about is "situation". Unfortunately, it's already in use as a technical term in another theory, where it's used to mean something that I don't mean here. (Barwise and Perry's situation semantics – see e.g. Barwise and Perry 1981.) I'll try, for the most part, to avoid "situation" in favor of "position" or "predicament". But please bear in mind that, when I *do* use "situation", I mean it to be interchangeable with these – I'm not using it in Barwise and Perry's technical sense.

I will, for now, just work unquestioningly with this sort of *positional* notion of context, and look at an array of examples that make trouble for the view that it's only the (positional) context of the *speaker* that's relevant to determining the proposition that a given utterance conveys (semantically) to its audience. I'll cast the argument as an argument that, since the use of a sentence in a single speaker-context can express different propositions to different audience members, we need to include a role for the positional context of the audience member, as well as that of the speaker, in the assignment of semantic values to context-sensitive vocabulary.

(A note about what is, for our purposes here, a notational choice: Lewis uses <world, time, location> triples as his contexts. I will talk about <world, time, individual> triples instead in what follows, since I think it makes the issues a bit clearer. For present purposes, nothing much will hang on the difference. Either will, in the cases we'll be concerned with, do the same work in fixing the rest of the features that we care about.)⁴

Later, we'll look at ways to take the arguments that follow as arguments against positional notions of context, and in favor of some other, less speaker-centric notion. (Such as the *presuppositional* notion of context that we find in Stalnaker (1979, 1998, and elsewhere) and the picture of context as *conversational score* that it's natural to take from Lewis (1979a). It's a bit surprising to find such an argument – on the face of it, it looks as if

⁴ Both run into trouble in other cases. The <world, time, location> line has trouble with cases of spatially overlapping speakers. Two interpenetrating ghosts simultaneously say "I am the ghost of Jacob Marley!" One speaks the truth while the other asserts a falsehood, since their uses of "I" pick out different individuals. The <world, time, individual> line has trouble with cases in which a time traveler loops back on himself. (I owe this unhappy discovery to Shen-yi Liao (MS).) Marty's younger self points to Fred, his older self points to Ed, and both say, "he is the killer!" Again, one speaks the truth while the other asserts a falsehood, since their uses of "he" pick out different individuals. For these reasons I think that it's probably best to treat *predicaments* – positions for individuals to occupy within worlds – as primitive, rather than attempting to identify them with either of these sorts of set-theoretic entities. (Another option is to identify them with time-slices of possible worldbound individuals, using whatever notion of a time-slice we need in order to be able to say that our time traveler has two distinct simultaneous time-slices. But this presupposes a lot of contentious metaphysics.) But for our purposes here, it's probably safe to ignore time travelers and interpenetrating ghosts, and think of predicaments as (happily represented by) whichever sort of triple we find most congenial.

sticking with a positional notion of context couldn't possibly do any harm, since these alternative notions can easily be brought under the positionalist's umbrella. Fix the world, time, and speaker of the utterance, and you'll have fixed which conversation the utterance is a part of, and the stage of the conversation at which the utterance occurs. You'll thereby have fixed the presuppositional context in which the utterance occurs. Indeed, you'll thereby have fixed *all* of the features of the conversational score in the conversation that the utterance is a part of at the time at which the utterance takes place. Fix the positional context, it's natural to think, and you get the presuppositional and scorekeeping contexts for free. So any work you could do with a presuppositional or scorekeeping context, you can do just as well with a positional context. This attractive thought turns out to be wrong, for reasons that we'll come back to later, but it is certainly an attractive thought.⁵

What's a *content*, exactly? This is a large and difficult question. It's difficult in part because it's hard to know which of the fights about "content" are terminological and which are substantive. One often has the feeling in these disputes that the parties to them are fighting about which labels to apply to various bits of a substantively common account of what's actually happening in communication. What I will do here is try to identify a particular notion of content which I take to be particularly interesting and important, and argue for some conclusions about what *that* sort of content depends on. The sort of content that I'll be concerned with is a sort that aims to capture *representational import* – what we want from a content (in this sense) is something that tells us how a given representational item represents things as being. The representational import of an *utterance* is tightly bound up with its communicative upshot. For an attributed content to capture how an utterance represents, it must capture how members of the utterance's audience must take things to be in

⁵ See Lewis 1980 p. 28 for something that looks a bit like this argument.

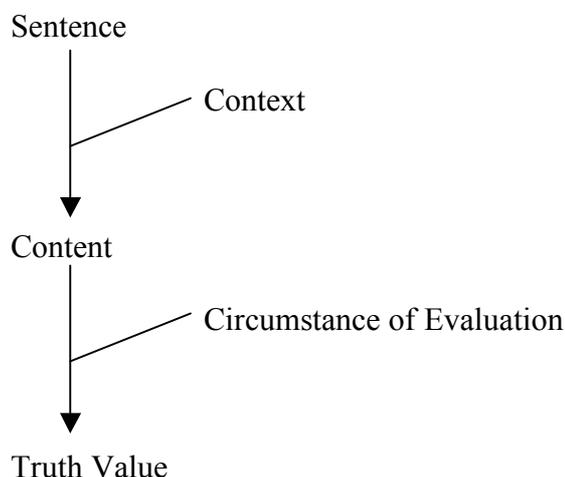
order to accept the utterance. To use a phrase stolen from Thony Gillies, the content of an utterance is “what travels” – it’s the message that the speaker conveys to her audience by means of the utterance.

Of course, the semantic content of an utterance isn’t *all* that’s conveyed to the audience. All sorts of other things besides the semantic content are conveyed via pragmatic means. I’ll be assuming a broadly Gricean picture, on which there’s a central semantic content that’s conveyed in virtue of the speaker’s and hearer’s fairly narrowly-drawn linguistic competence, derived from the language’s compositional semantics, and on which the rest of the communicative upshot of the utterance is parasitic, and derived in multifarious and largely non-compositional ways. (A familiar example: When you assert, “O’Leary has excellent handwriting” you communicate to me both *that O’Leary has excellent handwriting* and *that O’Leary is a lousy philosopher*. Your communication of the second thing is parasitic on your communication of the first.) The picture here is the familiar one on which one’s linguistic competence delivers a compositionally derived content, and then the fact that *that* content was delivered, plus a lot of potentially not-at-all-linguistic information, gets leveraged into a bunch of extra messages conveyed. This is, no doubt, a picture that will need tweaking and refinement in any number of ways, but it should be enough to be getting on with for now. That’s the sort of utterance content that I’ll be concerned with in what follows – the sort of content that captures an utterance’s *central, compositionally derived take-home message*.⁶

So much for ground-clearing. On to the plan for the remainder of the paper.

⁶ There are other roles we’ll want content to play – for example, governing ascriptions of truth-values, playing well in embeddings, etc. But we’ll certainly want something to play this role, and it seems quite plausible that what plays this role will be well-suited to play the others. It’s plausible to think that it’s the representational import of sentences – or anyway, something *very* closely connected with their representational import – that determines how they contribute to embeddings, for example.

The Kaplanian way of thinking about context-dependence described at the beginning of the section, and illustrated with the following diagram, is nowadays completely standard:



Post-Kaplan, it's either the orthodoxy, or as close to orthodoxy as anything we're likely to find in the philosophy of language, that *some* version of this picture is correct. It's riskier to claim that it's completely standard to accept the interpretation of the above diagram on which contexts are taken to be *speaker's positional contexts* and contents are understood in terms of central, compositionally derived, communicative import. But whether such a view is standard or not, it's at least quite widespread. I certainly used to believe it, almost nobody ever complained when I asserted or presupposed it, and people (well, philosophers at least) quite often spontaneously produce it when asked to give an account of how context-dependence works and pressed to explain the notions of "context" and "content".⁷

In what follows, I'm going to argue that this view of context dependence – on which it's always speaker's positional contexts that feed in to character in order to determine the

⁷ In the remainder of the paper, I'll consider at a few points what happens when we move to a different notion of context. I won't consider any rival notions of content, so this is the last time I'll mention the possibility of altering our interpretation of the Kaplanian picture of context-dependence by altering the role that we're taking *contents* to play.

content of an utterance – is mistaken. I'll do this by drawing attention to a number of examples about which it makes bad predictions. Before we start, let me try to identify, in minimally technical vocabulary, the consequences of this version of the Kaplanian picture – call it the *speaker-positional* theory of context-dependence – that I take to be mistaken.

One is the obvious fact that on this picture it's the context of *utterance* – understood as something that specifies the situation of the *speaker* – that's doing all the work in the fixing of content. So I'll be complaining about the speaker-positional theory's commitment to:

SPEAKER ONLY: Once we've fixed which sentence was uttered, content depends only on features of the speaker's (positional) context. Context-sensitive vocabulary is only sensitive to is features of the speaker's situation at the time of utterance.

I will argue in what follows that SPEAKER ONLY is false, and that we need a theory of context-dependence that allows for content to depend not just on the utterance's *origin*, but also on its *destination*. There are cases in which a single utterance semantically conveys different propositions to different members of its audience, so what a sentence conveys depends not just on the context in which it is uttered, but also on the context in which it is received. That is, I'll be arguing for:

AUDIENCE SENSITIVITY: For some uses of context-sensitive vocabulary, the contribution that they make to the content of sentences in which they occur is sensitive not (merely) to features of the speaker's predicament, but (also) to features of the predicaments of particular audience members.

(We need two different principles because, as will be illustrated below, AUDIENCE SENSITIVITY isn't just equivalent to the negation of SPEAKER ONLY.)

1. I'm Not Here Now

We'll start by looking at a couple of arguments against SPEAKER ONLY based on the communicative role of answering machine messages. Though I think that these arguments are ultimately unsuccessful, they'll serve to get some of the phenomena we'll be interested in on the table, and they'll point us toward the arguments that I think *do* succeed. (The following discussion is indebted, at too many points to note each of them separately, to Sidelle 1991 and Predelli 1998a.)

Consider the following case:

ANSWERING MACHINE

On Monday, Boris stands in his flat and records his answering machine message, by saying to his answering machine (which he's named Natasha): "I am not here now." He only makes the one utterance. Call Boris's utterance B, and call the context in which Boris makes his utterance PLANT. On Tuesday, Rocky calls, and Natasha plays back Boris's recorded message. Call the context of this replaying D1. On Wednesday, Bullwinkle calls, and Natasha plays back the message again. Call the context of this second replaying D2.

The standard way of thinking about context-sensitivity says: used in a context *c*, "I am here now" expresses the proposition that you get by composing the semantic values that the

component expressions receive in c – a proposition of the form, $[I_c \text{ am not here}_c \text{ now}_c]$. So Boris's utterance on Monday expresses the proposition, $[I_{\text{PLANT}} \text{ am not here}_{\text{PLANT}} \text{ now}_{\text{PLANT}}]$ – that is (assuming a natural view about the characters of “I”, “here”, and “now”) *that Boris is not in Boris's flat on Monday*.

It would be bad, obviously, if we had to say that the replayings of Boris's message just present their audiences with Boris's original utterance again, with the same content that it expressed in its original context of utterance. What's expressed to Rocky when he calls on Tuesday and hears the message pretty clearly *isn't* the proposition *that Boris is not in Boris's flat on Monday*. For one thing, that's just pretty obviously the wrong content to be expressed - it's about the wrong time, and it's false. For another thing, even if that *were* what was expressed, it is incredibly unlikely that Rocky would *know* that that was what was expressed, since he's unlikely to know just when Boris recorded his message.

Since Boris ought to be able to use his answering machine to transparently convey something true about his location on Tuesday to his callers on Tuesday, rather than to express, in some way his audience has no hope of picking up on, some false proposition about his location on Monday, we want a theory that tells us that what's conveyed to Rocky is the proposition, *that Boris is not in Boris's flat on Tuesday*. (And that what's conveyed to Bullwinkle is *that Boris is not in Boris's flat on Wednesday*.)

One way to deliver this result is to say that Boris's single utterance B has different contents on each of its replayings, and so that its content is sensitive not just to the context in which it's uttered, but also the context in which it's replayed, or perhaps the context in which it's heard. If this were the *only* way to deliver the result, then we'd have an argument against SPEAKER ONLY. This is an account according to which a single utterance B, which happens in just a single context of utterance, receives a different content relative to two different

replayings. In order to allow for this, we need to say that the context of the thing doing the replaying, or perhaps of the person hearing the replaying, and *not* just the context of the speaker, plays a role in determining the content of B. (Note that this is one of the places where denying SPEAKER ONLY and accepting AUDIENCE SENSITIVITY can come apart.)

But this isn't the only way to deliver the right result about what's expressed to Rocky and Bullwinkle. We can also say – and say quite plausibly, I think – that the right way to think about answering machines is to think of them as *utterance bombs*: What happens when Boris makes his recording by producing B is to plant an utterance bomb, which is later detonated when Rocky calls on Tuesday to produce a new utterance of “I am not here now.” The content of this newly-produced utterance is sensitive only to the relevant features of D1 – the context in which the replaying takes place. In the relevant sense of “speaker”, we should say that the speaker on each of the replayings is *Natasha*, not Boris, and that the context which is relevant to determining the semantic values of the context-sensitive vocabulary on each replaying is Natasha's context at the time of replaying. This lets us keep SPEAKER ONLY by revising our view about who or what counts as the speaker for purposes of determining which proposition is expressed when the message is replayed.⁸

There's an obvious complication about this view, which might lead one to think that adopting an utterance-bomb view won't actually be of any help in allowing us to retain SPEAKER ONLY. The utterance-bomb account clearly lets us assign the intuitively correct semantic value to “now” on each of the replayings. It's less clear that it lets us assign the correct semantic value to “I”. On the face of it, if we resolve all of the context-sensitivity in “I am not here now” with D1 – if we say that the proposition expressed to Rocky when

⁸ This is a view in the same spirit as what Sidelle (1991) calls *deferred utterance*, and Predelli (1998b) calls the *Remote Utterance View*. One difference which is at first glance very substantial, but I suspect is ultimately terminological, is that Sidelle is not willing to say that the answering machine is the speaker (for reasons that we're about to get to).

Boris's message is replayed on Tuesday is $[I_{D1} \text{ am not here}_{D1} \text{ now}_{D1}]$ – we'll have to say that "I" refers to the speaker of D1 (that is, to Natasha) rather than to Boris, the speaker of PLANT. In that case, the replaying that Rocky hears will express the proposition *that Natasha is not in Boris's flat on Tuesday*. And that proposition, while it's about the right *time*, isn't about the right *individual*.

One way to fix this second difficulty, and deliver the result that Natasha's replaying of the message on Tuesday expresses the proposition *that Boris is not in Boris's flat on Tuesday*, is to allow for a sort of *context mixing*, and say that some of the context-sensitive expressions in the replayed sentence have their semantic values fixed by the context of replaying, while others have their semantic values fixed by the context of planting. (At this point, it probably stops mattering very much which of the context of planting and the context of detonation we decide to call "the context of utterance", and which of Boris and Natasha we call "the speaker".) We could say that the proposition expressed by Natasha's replaying in D1 is not $[I_{D1} \text{ am not here}_{D1} \text{ now}_{D1}]$, but is rather $[I_{PLANT} \text{ am not here}_{D1} \text{ now}_{D1}]$ (or maybe $[I_{PLANT} \text{ am not here}_{PLANT} \text{ now}_{D1}]$). This would again require us to give up SPEAKER ONLY – whichever of Boris or Natasha we think of as the speaker, the content of the replaying is sensitive to features of some context other than theirs. (Though not the context of any audience member – this is another instance of giving up SPEAKER ONLY without endorsing AUDIENCE SENSITIVITY.)⁹

If this were our only option, then we'd once again be committed to denying SPEAKER ONLY. But again, we have another option. We can instead stick to the claim that the content

⁹ Another alternative: perhaps we say that characters always take $\langle \text{context of planting, context of detonation} \rangle$ pairs as arguments: $[I_{\langle PLANT, D1 \rangle} \text{ am not here}_{\langle PLANT, D1 \rangle} \text{ now}_{\langle PLANT, D1 \rangle}]$. This effectively winds up treating every utterance as a detonation of an utterance bomb – the difference between face-to-face conversations and answering machine messages then is just that in the first case, there's no distance between the context of planting and the context of detonation. Face to face utterances are, on such an account, just utterance bombs with very short fuses.

of the replaying is fixed entirely by features of Natasha's context at the time of replaying, and conclude from the example that we need to revise our first-pass, knee-jerk theory of the character of "I". We don't need to say that the character of "I" takes every <world, time, individual> triple to its *individual* constituent. (In a slightly different vocabulary: we don't need to say that it takes every predicament to the individual whose predicament it is.)

A natural thing to think about the replaying of Boris's message in D1 is that what's happening is not that *Natasha* is addressing Rocky, but that *Boris* is addressing Rocky *through* Natasha. Though Natasha is making the sounds, and so in some interesting sense counts as the speaker, it still seems right to say that Boris is the *author* of, the original source of, the utterance. (I take it that this is at least part of the idea behind Sidelle's (1991) notion of deferred utterance.) So we can say that the right story about the character of "I" is that it takes a <world, time, individual> triple or predicament to *the author* of the utterance that's taking place there. (The notion of *being the author of* is, obviously, going to need some spelling out, but there does seem to be some suitable notion around to *be* spelled out, and which ought to be able, once spelled out, to bear the required theoretical load.)¹⁰

This forces us to abandon the first-pass, knee-jerk theory of the character of "I", but it's not clear how big a cost that really is. After all, we already knew that the analogous theory of the character of "here" was false – we certainly need a theory of the character of "here" that doesn't just deliver the location of utterance (or even a location that includes the location of utterance), in order to handle cases in which, for example, the backcountry guide

¹⁰ One nice feature of this story, incidentally, is that it seems to get the difference between the answering machine case and the case in which Boris builds a self-monitoring function into Natasha and records, in his own voice, an "I am malfunctioning" message, to be played when Natasha detects that she is not working properly. In *this* case, we want to say that what's being expressed is *that Natasha is malfunctioning*, not *that Boris is malfunctioning*. The account of the character of "I" in terms of utterance authorship promises to be able to get both this case and the answering machine case right. If the reporting apparatus is responsive to the self-monitoring apparatus in the appropriate way, the replaying of Boris's recorded "I am malfunctioning" message does seem to trace its origins back to Natasha in the right sort of way for her to count as its author.

points at a spot on the map and says, “we need to go here”. So this sort of departure from the first thing we might be inclined to say about the character of an indexical expression is by no means unprecedented.

Summing up: The answering machine based arguments against SPEAKER ONLY don’t succeed. If we adopt an utterance bomb view, and accept a more sophisticated theory of the character of “I”, we can retain SPEAKER ONLY while making all of the right predictions about the communicative import of answering machine messages.

There are other cases, though, give rise to similar arguments which cannot be defused in the same ways.

2. *Jesus Loves You*

Consider the following case:

BILLBOARD

Horton produces a billboard on which is written the sentence, “Jesus loves you”. Call the context in which this occurs INSCRIBE. Frank and Daniel each drive past the billboard and read it.

What is the message that the billboard conveys to Frank? What does Frank need to come to accept in order to take on board what the billboard says? On *any* kind of standard contextualist view, what’s expressed to a reader of the billboard will be [Jesus loves you_c], for some c. If c is INSCRIBE, “you” will, presumably, refer to some group, and the proposition expressed will be something of the form, *Jesus loves G*, where G is the group that Horton had in mind at the time of his writing, or the group containing all of the people who

will ever read the billboard, or everyone who drives down highway 151, or... But that seems wrong. The overwhelmingly natural thing to say is that what the billboard conveys to Frank is *that Jesus loves Frank*.¹¹ In general, we ought to say that the billboard expresses, to each reader, the relevant singular proposition about *them*. (To reinforce this intuition, think about the effect of followup billboards a bit further down the road that read “I mean you!”, or “that’s right, buddy – *you!*”.) In order to deliver this result, though, we need to give up SPEAKER ONLY and accept AUDIENCE SENSITIVITY.

An initially attractive response is to say that billboards, like answering machines, are utterance bombs. We can then say that it’s the context of the billboard at the time of reading that fixes the semantic values of the context-sensitive vocabulary on the billboard. So when Frank reads the billboard, “you” picks out Frank, and when Daniel reads the billboard a bit later, “you” picks out Daniel.

I think that in fact we *ought* to say that billboards are utterance bombs, but that saying this doesn’t help us to retain SPEAKER ONLY. It doesn’t help us retain SPEAKER ONLY because in the case of billboards, we’ve got the possibility of multiple simultaneous audience members.¹² In the version of our case in which Frank and Daniel read the billboard at precisely the same time, the utterance bomb move is no help at all. The context of *the billboard* at the time of the two simultaneous readings won’t help us deliver different contents to Frank and to Daniel. If the character of “you” is a function that takes a context of utterance (which is, in this case, a context of detonation) as argument and delivers a group or individual as semantic value, we’re only going to get one referent per context of utterance.

¹¹ What about the self-locating proposition that’s true of all and only the objects of Jesus’s love? I don’t think that this is a good candidate, for reasons that we’ll get to at the end of the paper.

¹² Depending on the sort of technology in use, we’ve probably got the same possibility for answering machines, or at least for voicemail, as well. The billboard example is nicer to work with, though, since it’s a particularly clear case of this sort of phenomenon, and it avoids some distracting complications of the voicemail case.

And we've only got one (positional) context of the billboard to work with, since Frank and Daniel are reading simultaneously. (Thinking in terms of centered worlds, the only speaker's-end positional context that we have available to plug in to the character of "you" is <@, now, the billboard>. Other ways of modeling speaker positions do no better – they won't help us find more than one distinct predicament for the billboard to occupy at the time of detonation.)¹³ The fact that we've got to have the same input to the function that determines the content expressed to both Frank and Daniel prevents us from saying that the billboard expresses different propositions to each, and so it prevents us from saying the intuitively plausible thing, which is that the billboard expresses, to each reader, the singular proposition about them.

Moral: No theory of context-sensitivity which is only allowed to appeal to speaker's-end positional contexts will be able to deliver what seem to be the correct predictions about what Horton's billboard expresses to its readers. This is so regardless of whether we take the speaker (i.e., the originator of the relevant utterance) to be Horton, or take it to be the billboard itself. No specification of the properties of the originator of the utterance (whoever or whatever we take the originator of the utterance to be) is going to be of any help. What we need is sensitivity to the context of the *reader*. In order to handle these cases correctly, we need to abandon SPEAKER ONLY and endorse AUDIENCE SENSITIVITY.

¹³ We should resist the temptation to think that using <world, time, location> triples rather than <world, time, individual> triples could be of some help here, by allowing us to take the locations in question to be something other than the narrowly-drawn location occupied by the speaker. Making the locations *smaller* than the speaker will do no good – there's no prospect of finding semantically relevant differences between the various subregions of the region exactly occupied by the speaker, and no prospect of finding a principled way to claim that one of these subregions is the one relevant to fixing what's expressed to Frank, while a different one is relevant to fixing what's expressed to Daniel. If we make the locations *larger* than the speaker, we're no longer working with a notion of context that's *positional* in the intended sense. What makes a theory positional in the relevant sense isn't just that its contexts make reference to some position, but that they be well-suited to capture the thought that a context is a potential point of origin for an utterance. Later in the paper, we'll discuss responses which involve giving up the speaker-positional theory by abandoning a positional notion of context – I take this to be a species of such a response. (Thanks to Carrie Jenkins for discussion here.)

Note the parallel with the answering machine case. There, we only had the one context of original utterance, and two contents we wanted to have expressed. So we needed to add some dependency on a context other than the one in which the original planting utterance was made. The solution there was to say that the context of detonation – that is, *the context of the answering machine* at the time of detonation – is what fixes the semantic values of the context-sensitive expressions. Now, we've got the same sort of problem: We've got just the one context of detonation, but two contents we want to have expressed. To account for the way in which the proposition expressed by the billboard varies across different readers, we need to add some dependency on a context other than the one the billboard's in at the time of reading. In this case, that means adding some dependence on context of the *audience member*, not just context of *speaker*. Which means giving up SPEAKER ONLY and signing up for AUDIENCE SENSITIVITY.

3. *Televangelists, Tony Robbins, and Related Phenomena*

At this stage, it will be helpful to pause to address some potential objections and concerns, look at some more examples of the same sort of phenomenon, and get clearer about the argument against SPEAKER ONLY and in favor of AUDIENCE SENSITIVITY.

The first potential concern is that perhaps this is just a peculiarity of written language, which we don't need to take too seriously. It's not. In order to convey his message(s) to the world, Horton needn't make a billboard. He could achieve the same effect – conveying to each of a large number of people the relevant singular proposition about them – by addressing a large crowd, either on television or in person. (We'll focus on the studio audience, to avoid the appearance of any of our conclusions hinging on peculiarities of broadcast media.)

Horton, in the course of a sermon, says “Jesus loves you”. The natural way to hear this is as being relevantly like the billboard case. What’s expressed to Daniel is the singular proposition about Daniel, and what’s expressed to Frank is the singular proposition about Frank. It doesn’t seem to correctly capture the communicative import of Horton’s utterance to say that it expresses, to every member of the audience, the general proposition about everybody, or everybody in the audience, or all of the faithful, or... (That is, some proposition of the form, *Loves(Jesus, G)* for some group *G*. Call these *group propositions*.) And if we want to endorse a view on which Horton expresses different singular propositions to different members of his audience, rather than expressing a single proposition about the group to everyone, we’re going to need to give up SPEAKER ONLY and accept AUDIENCE SENSITIVITY. What seems to be happening here is what we might think of as a sort of *shotgun assertion*, in which different asserted contents are going out to different audience members, rather than a single content going out to all of them. Each audience member gets their own assertion-pellet, loaded with its own proprietary content.

(Actually, there are probably sorts of sermons for which different readings are appropriate – there certainly are communities and contexts in which the natural way to hear an utterance of “Jesus loves you” is as expressing, to everybody, the proposition about the group. What’s interesting is that there are also contexts in which that seems wrong – in which the natural way to hear the utterance is as expressing, to each member of the audience, the singular proposition about them in particular. (One way to bring out the difference is to think about, on the one hand, religious traditions in which what’s of crucial importance is the social, community-based features of collective worship, and on the other hand, traditions in which what’s of crucial importance is one’s particular, individual relationship with the divine.) A view that endorses SPEAKER ONLY can easily accommodate the first, group-

emphasizing sort of use, but seems to have trouble with accommodating the second, individual-emphasizing sort of use, and with explaining the availability of the two different ways of hearing utterances of “Jesus loves you”.)

It’s possible, of course, to resist this conclusion by denying the initial claim about what Horton’s utterance expresses to the various members of his audience, and what the billboard expresses to its readers. There are two proposals on the table as to the semantic contribution made by the occurrence of “you” in these cases. We can say that its contribution is always an *individual*, and so that the proposition expressed is always some singular proposition – something we can think of as being of the form, *Loves(Jesus,x)*. Alternatively, we can say that its contribution is always a group, and so the proposition expressed is always some group proposition – something we can think of as being of the form, *Loves(Jesus, G)*. If we say the first thing, we seem to be forced to give up SPEAKER ONLY and accept AUDIENCE SENSITIVITY – if any singular proposition is getting expressed to Frank, it had better be the one about *Frank*, and if any singular proposition is getting expressed to Daniel, it had better be the one about *Daniel*. If we say the second thing, we’re free to retain SPEAKER ONLY, by saying that the same group proposition is expressed to both Frank and Daniel.

A defender of SPEAKER ONLY could, at this point, dig in her heels and insist that the semantic contribution of the occurrence of “you” in Horton’s use of “Jesus loves you” really *is* always a group (for example, the group containing all and only people present for the sermon, or all of the members of Horton’s church), rather than an individual (for example, Frank), and that the intuitions to the contrary that I’ve been appealing to are not to be trusted. This sort of skepticism about the intuitions I’ve appealed to above can be bolstered by pointing out that, while the two proposals make different predictions about the *semantic*

content that Horton's uttered sentence has for its various listeners, in most cases they *won't* differ in their predictions about the net communicative upshot of Horton's utterance.

Suppose that what's expressed is the group proposition. If Frank knows that he's a member of the relevant group (as he typically will), he's in a position to work out very straightforwardly that if the group proposition is true, then the singular proposition about him must be true as well. So acceptance of the group proposition is very likely to bring acceptance of the singular proposition in its wake.

Suppose that what's expressed is the singular proposition. If Frank knows a bit about who else is being addressed, he'll know that Horton is aiming to assert similar singular propositions to each of the members of the relevant group. And if he takes Horton to be, in general, a reliable source of information about the distribution of Jesus's love, he's likely to accept that each of *those* assertions is correct as well. So acceptance of the singular proposition is likely to bring acceptance of the group proposition in its wake.¹⁴

¹⁴ There are other possible proposals as to what "you" contributes in these cases. One is that it makes *no* contribution, and so that what Horton expresses is not a proposition at all, but a *propositional radical* of the sort Kent Bach (1994 and elsewhere) discusses. Another possibility is that it makes some rather fancy sort of contribution, such that what's expressed by the sentence as a whole is a *self-locating* proposition (equivalently, a property - see Lewis 1979b, Chisholm 1982), or a *relative* proposition of the sort we find in MacFarlane (2003, 2005, forthcoming) and Lasersohn (2005). I've argued elsewhere for the self-locating version of this sort of view about other natural language sentences (see Egan 2006, 2007, forthcoming). Toward the end of the paper I'll say why I don't think that's the right response here. For now, let's set these aside - we'll return to them briefly in the final section.

One proposal that we *should* draw some attention to now is the proposal that what "you" contributes is not any sort of referent at all, but a variable, which is bound by an unpronounced quantifier. On this sort of proposal, what Horton expresses is a single *quantified* proposition - something we can think of as being of the form, $\lambda x[\text{Loves}(\text{Jesus}, x)]$. There is a plausible precedent for this sort of use of "you" - it's the use in which "you" is more or less interchangeable with "one". (As in, "You can't always get what you want", or "To get to the store, you go down three blocks and you turn left.") The best thing for the advocate of this sort of view to say is that the quantifier in question is a restricted universal quantifier, with a contextually determined restriction (to, for example, the people present at the sermon). This sort of proposal will yield the same possible-worlds truth conditions as a proposal on which "you" simply contributes the group that's restricting the quantifier. They'll be different only in the claims that they make about how the sentence manages to compositionally arrive at those truth conditions.

A few reasons why I don't give this proposal more time in the main text of the paper: First, its plausibility hinges on special features of "you" that aren't common to all of the troublemaking examples that we'll be looking at, and so it doesn't seem to be viable as a *general* defense of SPEAKER ONLY. (Quite a few of them admit of this sort of response, but not all of them do.) Second, while it does help with some of the

There is a lot more to say here about how to attempt to disentangle intuitions about net communicative upshot from intuitions about semantic content, and how to adjudicate questions about which parts of the net communicative upshot of some communicative event are central and which are parasitic. But I think we needn't take up those issues here. I do think that, at the end of the day, the direct intuition about what the billboard expresses to each of Frank and Daniel is pretty compelling – compelling enough that denying it really looks very unattractive. But we can do more to motivate the singular proposition view than just banging the table and saying “come on, doesn't that just seem right?” over and over.

Here is another case – a secular one, this time – of the same sort of phenomenon in spoken language: Tony Robbins says, in the course of a motivational seminar, “You can take control of your life!” Suppose Frank and Daniel are both in attendance at the seminar. Once again, the natural thing to say is that what's conveyed to Frank is the singular proposition about Frank, and what's conveyed to Daniel is the singular proposition about Daniel. A nice feature of this case for our purposes, is that Robbins uses the singular “life” rather than plural “lives.” This seems to force us to say that his uses of “you” are contributing an *individual* rather than a *group*. If his use of “you” was picking out a group rather than an individual, Robbins would have had to say “lives” rather than “life”.

arguments I'll be running, the fact that it delivers the same possible-worlds truth conditions as the proposal on which “you” contributes a group makes it, for the most part, vulnerable to the same complaints. Finally, there are some difficulties with assimilating Horton's use of “Jesus loves you” in the sermon to the sorts of “impersonal” uses of “you” that provide the precedent for the quantified proposal. The cases where we have a likely precedent for a use of “you” that contributes a variable that's bound by an unpronounced quantifier are cases where it looks like the quantifier is probably a *generic*, not a restricted universal. But using the generic gets the meaning *radically* wrong – $GEN(x)[Loves(Jesus,x)]$ is about as ill-suited for capturing the direct, “I mean you in particular” import of Horton's use of “Jesus loves you” as anything one could hope to find. And if the cases that provide a precedent for bringing in any sort of quantifier at all are ones in which the quantifier in question is a generic, it starts to look suspiciously *ad hoc* to say that *here* we find a restricted universal instead. (Though the semantics of “one”, and the associated uses of “you”, are still up in the air, so it's far from settled what the right story is going to be. See for example Moltmann forthcoming.) In general: the “impersonal” use of “you,” on which it's more or less interchangeable with “one”, that motivates the quantified proposal seems to be exactly the *wrong* place to look for an account of Horton's use of “Jesus loves you”. One would be hard pressed to find a way to do more violence to Horton's meaning than by replacing “you” with “one”.

Another potential source of arguments against the proposal that “you” is picking out some group in these cases is looking at the analogous sentences in languages that explicitly mark the difference between singular and plural second-person pronouns. I’m not really qualified to run these arguments properly, but the preliminary reports from my native-speaker informants about German, Dutch, Danish, and Norwegian suggest that, in those languages anyway, we can construct versions of Horton’s and Robbins’s sentences with unambiguously singular second-person pronouns, which it would be felicitous for them to use in the relevant sorts of circumstances. (There’s some dispute between my different German-speaking informants about the Robbins case, though they’re unanimous about the “Jesus loves you” case.)

Above, I very briefly attempted to reinforce the intuitions in favor of the singular proposition view by asking you to think about the effect of follow-up signs a bit down the road, saying things like, “I mean *you!*”, or “Yeah, that’s right – *you!*”. Another good case for motivating singular proposition view is the old Uncle Sam recruiting posters for the US Army, which featured a picture of Uncle Sam, pointing his finger out of the page, with the caption, “I want *you* for the US Army”. The focal stress helps in both cases. What’s being highlighted by the stress on *you* isn’t the special importance to Jesus of drivers on highway 151, or to the US Army of the group composed of all of the readers of a particular sign. (Or whatever. Messing with the story about just which group is getting picked out doesn’t seem likely to help here.) What’s being highlighted, to each reader, is *their own* special importance to either Jesus or the US Army. This is easy to explain on the singular proposition view, difficult on the group proposition view.

In general, views according to which “you” contributes a group rather than an individual have trouble with phenomena to do with *contrast*. Here are some things that

Horton could go on to say in his sermon, which are easier to make sense of if his initial utterance of “Jesus loves you” is understood as expressing different singular propositions to each member of the audience than if it’s understood as expressing the same proposition about the group to everyone.

Suppose Horton goes on to say either of the following:¹⁵

- 1) Don’t misunderstand – I don’t just mean that Jesus cares about everybody impersonally, or even everybody religious, or everybody who’s at this sermon. I mean that Jesus loves *you*.
- 2) Of course Jesus also loves the person sitting beside you – indeed, Jesus loves this entire congregation – but what it’s important to focus on, the crucially important fact, is that Jesus loves *you*!

These, and other easily-constructed sentences in the neighborhood, seem to be *contrasting* the claim expressed by “Jesus loves you” with the various propositions in the neighborhood that represent some *group* or other as being among the objects of Jesus’s love. That’s hard to make sense of if what’s expressed by “Jesus loves you” just *is* one of these group propositions. If “you” just contributes the group containing all and only the people at the sermon, there shouldn’t be any interesting contrast available between the content (in Horton’s context) of “Jesus loves you” and that of “Jesus loves the people at this sermon” that Horton could be making by saying, e.g., “What’s important isn’t that Jesus loves the people at this sermon – what’s important is that Jesus loves *you*”. If the two sentences just have the same content in Horton’s context, then there’s no contrast there to emphasize. (Note

¹⁵ Thanks to Hud Hudson for these.

that the contrastive claims continue to make sense however we explicitly fill in the group – so fancy footwork about just which group is getting picked out doesn't seem likely to help.)¹⁶

There are also some things that Tony Robbins could go on to say that apply more pressure toward the singular reading:

- 3) Even if you watch all of the rest of the people in this room fail – even if they all forget the lessons we've learned here and fail miserably in their lives – you can still succeed.
- 4) Never mind what the bozos sitting next to you do! You can take control of your life!

These are hard to make sense of – indeed, they seem likely to turn out to be self-contradictory – if “you” is contributing a group rather than an individual.¹⁷

¹⁶ There are cases in which one can felicitously contrast a sentence S1 with a sentence S2 that has the same truth conditions – for example, “what I want to emphasize here isn't that *Bill* kissed Mary, but that Bill *kissed* Mary”. In these cases, it looks as if what's being contrasted isn't the different contents of the two sentences, but two different classes of contrasts with the common content of both sentences. (What's important isn't the contrast between Bill's kissing Mary and, e.g., Joan's kissing Mary, but the contrast between Bill's kissing Mary and, e.g., Bill's kicking Mary.) This doesn't seem to be what's going on in the cases in question. For one thing, these sorts of contrasting-contrasts cases seem to require that the speaker stress elements of *both* of the clauses being used to make the contrast – “what I want to emphasize here isn't that Bill kissed Mary, but that Bill *kissed* Mary”, with no stress at all on any element of the first clause, sounds pretty bad. For another, it's a bit hard to see what the relevant pairs of contrasts would be in the cases in question. On the hypothesis that “you” is contributing an individual rather than a group, though, it's completely straightforward to tell a very natural-sounding story about what's being emphasized, and what's being contrasted with what.

¹⁷ There is a sense of oddness about (3) and (4), and there are some interesting cases in the neighborhood that are very clearly defective. I think that another advantage of the shotgun-assertion approach is that it lets us give a better diagnosis of the sort of oddity or defectiveness they display than the group-proposition approach (or the quantified-proposition approach). Suppose Robbins says, “Only you can take control of your life – the rest of these people are doomed.” There's certainly *something* wrong with this. On a group or quantified proposal, what's wrong with it is that it's (nearly) *contradictory*. It's true only if every member of the audience is such that they alone among the audience members can take control of their lives. On the shotgun-assertion proposal, what's wrong with it is a particularly transparent sort of *insincerity*. What Robbins asserts to Frank isn't impossible (or at least, incompatible with there being more than one person at the seminar), but it's incompatible with what he asserts to Daniel.

It's helpful to compare this version of the Robbins case to the cases of the deceptive crime boss and the overenthusiastic sportscaster: The crime boss hopes to get rid of his crew and pocket all the money after the job. So he pulls each member of the group aside and says to them, “You're the only member of the gang I trust. Once the job is over you and I will kill the others and split the proceeds. Wait for my signal – when I shoot the first one, help me with the rest.” The crime boss has made a bunch of assertions, one to each member of his crew, no two of which could both be true. (Thanks to Denis Robinson for this case.)

It's worth looking at some other instances of the same sort of phenomenon, to assuage concerns that that these are just isolated, one-off cases that we can safely ignore.

In the briefing room before a dangerous mission, the commanding officer says:

- 5) Look around you. There are 100 people in this briefing room. The mission you're about to undertake is extremely dangerous. It's possible that, two hours from now, you will be the sole survivor of this group.
- 6) You know your assignment.
- 7) For security reasons, it's very important that you not tell anyone else about your orders – not even the other people in this room.

Before the show, the game-show host says, to a group of would-be contestants:

- 8) Only one of you will be selected. But be ready - it's possible that you will be that person.

The alarmist newscaster warns:

- 9) The maniac is still on the loose! Your family might be in danger!

The sportscaster, in a bout of mad enthusiasm, says "Chauncey Billups is the greatest player of all time! And so is Rip Hamilton! And so is Rasheed Wallace! Heck, the Pistons are a team composed entirely of guys who are the greatest basketball player of all time!" The sportscaster's made a single assertion which is incompatible with the fact that the Pistons have more than one player on the team.

The Robbins case, I think, is more happily assimilated to the crime boss case than the sportscaster case. What he's done wrong isn't to say a single, crazily false thing to everybody. It's to commit himself to a package of claims (each made to different people) such that he can't sincerely and unconfusedly make all of them at once. In fact, it's probably best assimilated to a version of the crime boss case in which the boss doesn't bother to pull each member of the gang aside and talk to them separately, but instead just delivers the relevant speech to each member of the gang in sequence, in full view and within earshot of the other members of the gang. Here it's not just that the boss's package of utterances is insincere, but that it's insincere in a particularly transparent way, that makes it ill-suited for achieving its ends.

(This sort of case is reminiscent of the old Smokey the Bear ads, in which children watching Saturday morning cartoons were told, "only you can prevent forest fires". I suspect that I wasn't the only child that those ads left feeling a bit confused and overburdened.)

The producer says, to the collected contestants on the new season of *Survivor*:

10) You're on your own. None of these people will help you when the chips are down.

Al Gore, in a speech on climate change, says:

11) Even though you're just one person, your efforts can make a difference.¹⁸

So far, all of the cases have involved the second-personal pronoun "you". But the same sort of phenomenon does seem to occur elsewhere.

Outside on a clear, starry night, the astronomy teacher says to her class:

12) "Pick a star. That star might be host to an entire civilization, just as complex and rich as our own."

The Dean says, to the incoming medical school class:

13) "Look at the person to your left. That could be the person who cures cancer in 20 years."

(At gender-segregated medical schools, one can generate similar examples with "he" and "she".)

Here we have uses of demonstratives (and third-person pronouns) that seem to need to pick up on different referents relative to different audience members.

Another trend so far has been that all of the examples have been cases of one-sided group address – cases in which there's one designated speaker who's doing most of the

¹⁸ It may help to think of this in the midst of, e.g., "You may be thinking: 'but can my individual actions really do anything?' Don't worry: Even though you're just one person, your efforts can make a difference. Let me tell you how..."

talking, and an audience that's not saying very much. These are not the only places to find cases that make trouble for SPEAKER ONLY and motivate AUDIENCE SENSITIVITY, though.

Here is another family of cases:

The undercover cop infiltrating the bank heist ring says, both for the benefit of the gangsters in the room and his fellow officers in the surveillance van,

14) "Everything's going just as we planned."

Here the reference of "we" seems to be sensitive to the audience's context. He expresses to the gangsters that everything's going as he and the gangsters planned, and to the surveillance team that everything is going as he and the rest of the police planned.

Before their SCUBA wedding, Gordon the groom says, to both the bride and the divemaster,

15) "I'm ready,"

He communicates to his fiancé *that Gordon is ready to get married*, and to the divemaster *that Gordon is ready to dive*. He does this because "ready" makes a different contribution in his communication with his fiancé than in his communication with the divemaster.

Linda is on the phone to her husband, who's asking whether she'll be home for dinner. Meanwhile, the chair of her department is asking if she'll be at the department picnic. She says, into her phone but while nodding to the chair,

16) "I'll be there."

She communicates to her husband *that Linda will be at home (at dinnertime)*, and to the chair *that Linda will be at the park (on the day of the picnic)*. She does this because "there" makes a different contribution to her communication with husband than it does to her communication with the chair.

A similar case – Linda’s husband says “we need a bottle of red, don’t we?” and the chair says, “Smith is your pick of the job candidates, right?”. Linda responds to both with a single utterance of,

17) “That’s true.”

She communicates to her husband that it’s true that they need a bottle of red, and to the chair that it’s true that Smith is her pick of the job candidates. She does this because “that” makes different contributions to the content she expresses to her husband and to the content she expresses to the chair.¹⁹

Here is a different sort of case, hinging on “now”: Krypton is going to be destroyed by a galactic cataclysm that’s about to originate in their system. As they prepare to evacuate, they broadcast a series of warning signals ahead of the wave of destruction, which will reach any civilizations in the endangered area a few years ahead of the cataclysm. (The signals and the cataclysm are both moving at the speed of light, so the interval between the arrival of the signal and the arrival of the cataclysm remains constant.) The first message says, in part:

18) “Those who act now can still save themselves.”

(Happily, this story takes place in one of those galaxies in which it’s common knowledge that everybody in the galaxy speaks English.)

A subsequent message says,

19) “By now, it is too late.”

¹⁹ A similar class of cases shows us why we shouldn’t say that the audience’s context does *all* the work (this is a view that we might call AUDIENCE ONLY): What’s happening in these cases is that, since Linda’s involved in two conversations at once, she can a single sentence to make different moves in the different conversations to which she is a party. It could also happen that the other participants in both conversations simultaneously use the same sentence, each with a very different communicative upshot. One such case is when both the chair and Linda’s husband say, “I’ll be there too.” Another is when two different speakers simultaneously say to you, while pointing at different people, “that guy’s one to watch out for”. Given the possibility of multiple incoming utterances of the same sentence, with importantly different origins, we need to allow for a role for both the speaker’s and the audience member’s context in the assignment of semantic values to context-sensitive vocabulary.

And a final warning says,

20) “The wavefront will arrive ten seconds from now. Nine. Eight. Seven...”

If we want to say that the Kryptonian transmissions in 18-20 express propositions about particular times, and take truth values relative to worlds, rather than expressing tensed propositions that take truth values relative to <world, time> pairs, then it looks as if we’ll need to say that they express different propositions – because about different times – to the different (successively more distant) civilizations that receive them. (To the Alpha Centaurians, the transmission in (19) expresses the proposition that, by t , it is too late. To the more distant Earthlings, it expresses the proposition that, by $t+$, it is too late.)

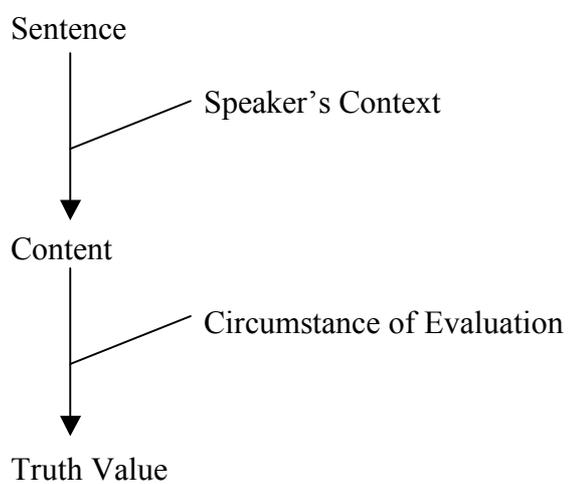
A single positional context on the speaker’s end can’t handle the sort of variation in what’s expressed to different audience members that we see in all of the above cases. To generate that sort of variation, you need an input to the character that captures the differences on the audience end.

The sort of sensitivity of semantic content to audience, as well as speaker, context that we saw in the billboard case isn’t just a special peculiarity of (certain sorts of) inscriptions. It turns up, in some pretty familiar and reasonably common cases, in spoken language as well. So the troublemaking phenomena aren’t just restricted to a few outlying “don’t care” cases, and it looks as if we really do have good reasons to give up the initially attractive picture of context-sensitivity that we began the paper with, on which the semantic values of context-sensitive expressions are sensitive only to the features of the context of the speaker (i.e., sensitive only to features of the speaker’s situation at the time of utterance).

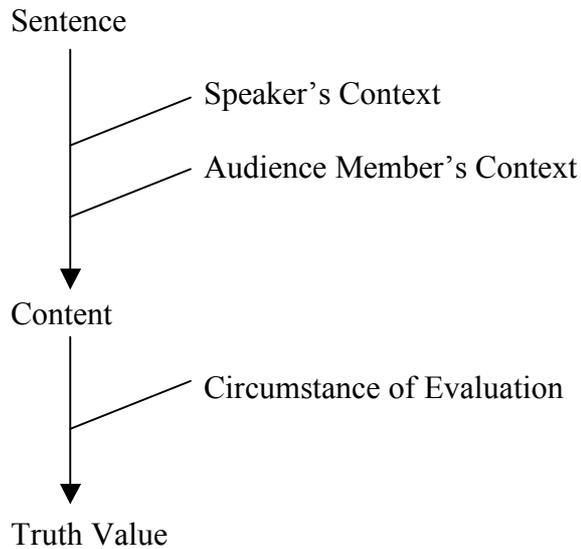
4. Content Proliferation vs. Utterance Proliferation

I've been arguing that we need to give up the picture according to which the semantic values of context-sensitive expressions are sensitive only to features of the speaker's (positional) context, and admit some sensitivity to features of the audience's (positional) context, as well.

Then we'll need to abandon the following sort of picture of how sentences determine truth values in contexts:



If we want to stick with a positional notion of context (we'll look at some other options in the next section) then what we'll need to endorse instead is some version of the following picture:



The main take-home message of this paper is that we ought to endorse the second picture rather than the first. There are, however, different sorts of views that fit the template provided by this second picture. I'll spend the last couple of sections exploring some different ways of spelling out the details of such a view.

One way to endorse the second picture is to say that, when Tony Robbins says “you can take control of your life”, he makes a single utterance that expresses different propositions to different members of his audience. Another is to say that Robbins makes a bunch of different utterances – one directed at each member of his audience – each of which has a unique content. One way or another, we need to give up the claim that Robbins makes a single utterance with a single content; we can do this either by allowing for a proliferation of *contents* (that is, by allowing for a surprisingly high ratio of contents to utterances) or by allowing for a proliferation of *utterances* (that is, by allowing for a surprisingly high ratio of utterances to physically-individuated productions of sounds, etc.).

It's not clear to me just how much hangs on which of these views we endorse. In some respects, this looks like a decision about which of the things that everyone ought to agree exists to attach the label "utterance" to. There are, however, a few considerations that seem to favor thinking about things in the *many utterances* way over thinking about them in a *content relativist* way, and I'll briefly canvass a few of them here.

It's not just context-sensitive vocabulary that one can use to express different propositions to different audiences. The con man says, "the money is in the bank", communicating to his dupe the false proposition that the money has been deposited in the financial institution, and to his confederate that the money's been buried by the side of the river. We can use ellipsis to the same effect – when I'm involved in two simultaneous conversations, I can use "I will" to communicate to one audience *that Egan will be at the party* and to another *that Egan will pay the cable bill on time this month*. In these cases, we haven't even got the same (syntactically individuated) *sentences* being directed at the different audience members. In the "bank" case, the con man's utterance to the dupe and to the confederate contain different lexical items, and in the "I will" case, the utterances I make to my two audiences contain different bits of elided material. In these cases, it's much more natural to say that the speaker's made one utterance to one audience and a different utterance to another audience, than to say that the speaker's single utterance has got different constituent structures relative to different audiences. And, thinking of utterances as something like *conversational moves*, I find it much more attractive to say that the speaker's made two distinct conversational moves, rather than making a single move that has very different upshots for different audiences. And once we're multiplying utterances in the ambiguity and ellipsis cases, it seems fairly natural to do so in the context-dependence cases as well.

Looking at other sorts of speech acts than assertion also bolsters the case for an utterance-proliferating interpretation of what's going on, I think. (Though again, I don't think any of these considerations are decisive.)

First consider *performatives*. Horton is performing a group wedding ceremony for 100 couples. At the crucial moment, he says, "I pronounce you husband and wife". The natural way to describe what happens here is, I think, that Horton has simultaneously performed 100 different performatives, one directed at each couple. This is the *shotgun wedding* picture – what Horton has done is fired a bunch of wedding pellets out into the crowd, and each couple is on the receiving end of one pellet. Then when Horton goes on to say "you may kiss the bride", he fires a follow-on barrage of permission pellets, rather than delivering a single blanket permission to the group as a whole.

Commands look similar. When grandma says "eat your oatmeal" to the collected group of grandkids, the natural thing to think is that she's delivered different commands to the different grandchildren. Each has been commanded to eat their own portion of oatmeal, and the command that I've been issued will be fully complied with if I eat mine, regardless of what happens to everybody else's. There's another sort of command grandma could have issued – she could have issued a collective oatmeal-eating command to the group, enjoining us to collectively bring it about that each grandchild has eaten his or her own particular portion of oatmeal. (Or that, somehow or other, all of the oatmeal be eaten by grandchildren – maybe it's okay if Ryan eats everybody's.) If she does this, I won't automatically be off the hook once I've eaten my oatmeal – if Mary hasn't eaten hers, I'm still liable to criticism as a member of a group that's failed to comply with a command to perform a collective action.

In general, there are two ways to hear commands directed at groups: One is a genuine group command, where I'm commanding the group to undertake some collective action, and there's collective responsibility to bring it about. The other is a case where everybody emerges with their own individual obligation, there isn't any collective obligation, and the individuals haven't dropped the command-following ball if they don't take steps to make sure the others comply with their obligations. This is, I think, best explained by saying that one can use, e.g., "eat your oatmeal", "do your chores", etc., either in a "blanket" way, to issue a single collective command to the group as a whole, or in a "shotgun" way, to issue many different commands that can be complied with individually, rather than single group command. And again, this would provide a precedent for the multiplication of speech acts to which, once we've got it in place, we might as well assimilate the phenomena that motivate AUDIENCE SENSITIVITY.

This isn't conclusive – we are free to say that there's a single command issued to the group, the upshot of which is just a bunch of narrowly individual obligations. But that doesn't really seem like the most perspicuous way to carve up the theoretical landscape. And really, I suspect that the best one can hope for as far as arguments for either content proliferation or utterance proliferation is this sort of consideration about the most perspicuous ways to carve up the territory.

5. Competing Notions of Context

All of this has proceeded on the assumption that the right notion of context is a *positional* notion. It's now time to revisit that assumption.

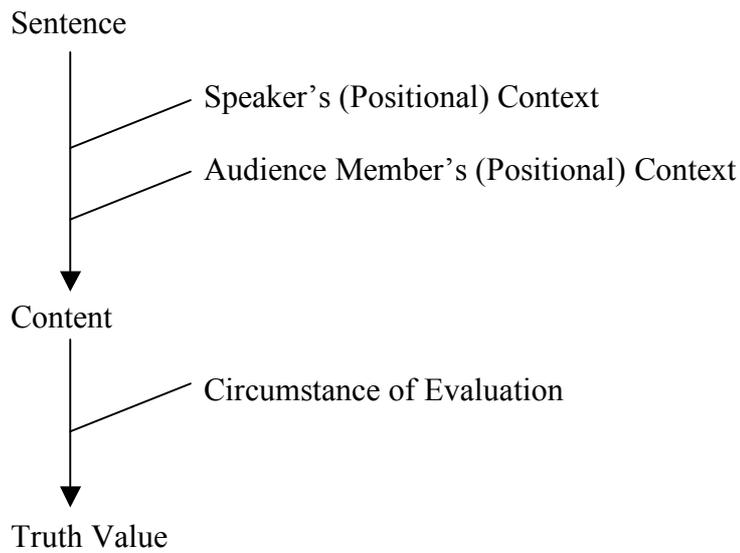
In the introduction, I offered a brief argument that a positional notion can do everything that a presuppositional or conversational-score notion of context can do, since

once you've fixed the speaker's position (for example, by picking out a <world, time, individual> triple), you'll thereby have fixed which conversation the utterance is a part of, as well as all of the potentially semantically relevant features of the current state of the conversation. Now we're in a position to see clearly what's wrong with that argument.

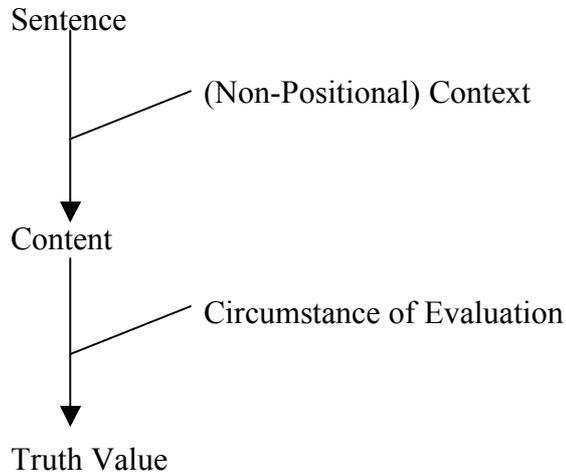
One thing that's wrong with it is that it ignores the possibility of cases in which a single speaker is simultaneously participating in two or more different conversations, each of which is in a relevantly different state. It ignores the possibility of using the same pronounced bunch of words to make different moves in different conversations. The cases in which Linda is simultaneously talking to her husband and her department chair are clear cases in which this is what's happening. In other cases, it's not clear that there are really two distinct *conversations* – the question of how to individuate conversations is going to be hard, and I'd prefer not to get bogged down in it just now – but it is clear that something relevantly similar is going on.²⁰

This leads us to another choice point. One option is to stick with a positional notion of content, and adopt the picture advocated above:

²⁰ Perhaps the speakers are involved in a single conversation in which different groups, both of which include the speaker, share relevantly different presuppositions, or perhaps different subgroups of the participants to a conversation's communication with each other is governed by different conversational scoreboards. The issue of whether the borders of conversations line up with the borders of the groups whose interactions are governed by different conversational scoreboards or not strikes me as largely terminological.



Another is to adopt a different, less speaker-centric notion of context, and retain a single-context picture:



Adopting the second sort of approach is in one respect easier – all we need to do is revise our view of what sort of thing feeds into the character of context-sensitive expressions. We can accommodate the troubling phenomena above by noting that, once we've got the right (presuppositional or conversational-score based) notion of context, which thing you

need to feed in to character to determine what's expressed to a given audience member by a use of a given sentence depends on more than just features of the speaker, and you might need to feed in different things in order to get the content expressed to different members of audience.²¹

There are, however, some phenomena that are more perspicuously represented by taking the first route – we can, for example, more easily isolate the different contributions made by the properties of the audience's situation and those made by the properties of the speaker's on the first approach than the second.

Perhaps we can get the benefits of both. Let me close with a proposal that attempts to do that. I'll first present the idea quite metaphorically, and then cash it out formally.

Forget everything we've just done for a moment. Assume a positional notion of context, and think of characters as functions from (positional) contexts to semantic values. Now consider the possibility that, sometimes, a context doesn't succeed in determining a semantic value for some expression – consider, that is, the possibility that characters are only *partial* functions from contexts to semantic values. In these cases, the final products of the compositional semantics would be *incomplete* – they'd be, to be even more metaphorical, would-be propositions with blank spots in them.²²

Now think about the natural story about the communicative import of such incomplete would-be propositions. What will happen with them, if they're to have any

²¹ This is also the sort of response that Predelli (1998b) advocates in response to a similar class of phenomena. It's worth noting that one could implement this sort of proposal by using <world, time, location> triples where the locations are big enough to include the locations of both the speaker and an audience member. This would be a theory of context that was in some sense a *positional* notion – it's picking out a location, after all – but not in the sense that I've been criticizing. Making *these* sorts of locations – ones that are bigger than the speaker's narrowly-drawn location – the inputs to character clearly isn't a way of making the resolution of context-sensitivity dependent only on the properties of the speaker at the time of utterance. (Since in determining the content expressed to Frank, we'll need to input a different big location than we use when determining the content expressed to Daniel.) Thanks to Carrie Jenkins for discussion here.

²² This is either a lot like Kent Bach's (1994, etc.) view about *propositional radicals*, or else it just is that view stated in slightly different terms, and motivated by quite different phenomena.

communicative import at all, is that each audience member will fill in the gaps in order to produce something that's a potential object of belief. But not just any filling-in is okay – the meaning of the expression, together with the original context of utterance, determines how the semantic value an audience member assigns to a particular gap in the would-be proposition depends on the features of the audience member's own particular predicament.

How do we cash this out formally? The best way, I think, is not to say that the characters of context-sensitive expressions are (sometimes) partial functions from speaker's positional contexts to semantic values. Instead we say that they are functions that take a speaker's positional context as argument, and deliver sometimes a semantic value, and sometimes a function from *audience* positional contexts to semantic values. This is a picture on which, as it were, the speaker's context gets the "first shot" at determining a semantic value for a context-sensitive expression, but sometimes the speaker's context "passes", and kicks the job down to the audience's context.

This lets us isolate the different contributions of speaker and audience context, while still keeping a formal picture on which characters just take speaker's contexts as arguments. (They now sometimes deliver some unorthodox values, though.)²³

This is formally equivalent to a story on which characters are really functions from ordered pairs of a speaker's context and an audience member's context to semantic values. But I think that thinking of things in this way is less conducive to perspicuously display the different contributions of the speaker's and audience's contexts. (In particular, to displaying the contribution the speaker's context makes to the way in which semantic value is sensitive (or not) to the audience member's context. Of course all of that is extractable from the

²³ Alternatively and equivalently, we could say that the character always delivers a function from audience contexts to semantic values, but that it's often a constant function.

functions from <speaker context, audience context> pairs to semantic values, but it's not going to be as prominently on display.)

6. *Why Not (Self-Locating) Relativism?*

Maybe there's a way of using 'relativist' according to which this view ought to be classified as relativist. I'm sometimes attracted to calling the version of the view on which there's a single utterance that expresses different contents to different audience members "content relativism". (Indeed, this is what John Hawthorne, Brian Weatherson and I called such views in Egan, Hawthorne and Weatherson 2005.) Whether to attach the term to the view or not strikes me as not a particularly interesting issue. What's certainly true, though, is that it's a very different sort of view than the sort that Hawthorne, Weatherson and I wound up advocating in that paper, the sort I've since advocated elsewhere under the name 'relativism', and the sort that John MacFarlane (2003, 2005, forthcoming), Peter Laserson (2005), and others have advocated under the same label.

There are important differences between the sorts of views that different relativists propose, and not all of them agree that the others' views deserve the name. (The view I've argued for under the name 'relativism', for example, doesn't count as relativist by MacFarlane's criteria.) Before concluding, I'll say a bit about why I don't think that we should adopt a relativist view in response to these phenomena – at least, not a relativist view of the sort that I favor in other areas.

Elsewhere (Egan, Hawthorne and Weatherson 2005, Egan 2006, 2007, forthcoming), I've argued for a view according to which some sentences of natural language have *self-locating* propositions – functions from possible positions, locations, or predicaments (modeled with centered worlds - triples of a world, an individual, and a time) to truth-values

– as their semantic values. (Equivalently: that some sentences of natural language have *properties* as their semantic values. See e.g. Lewis 1979b, Chisholm 1979.)²⁴ Why not say the same thing here? Indeed, why isn't such a proposition a *better* candidate to be what's expressed to Frank by, for example, his reading of Horton's "Jesus loves you" billboard than the singular proposition about Frank? After all, isn't Frank's coming to self-attribute *being loved by Jesus* an even more reliable effect of Frank's credulously reading the billboard than his coming to believe *that Jesus loves Frank*? Even if Frank is confused about who he is – if he thinks that he's Napoleon, or Hume – he'll still come to self-attribute the property, but he probably won't come to believe the singular proposition about Frank.

The first reason why I'm not inclined to go self-locating here is that: (a) I'd like a unified theory of indexicals, so that "I" and "you" work in roughly the same way, and are distinguished by their characters, rather than just working in radically different ways; and (b) the self-locating story for sentences involving "I" is, given the notion of content we're working with, a clear disaster. One way to see this is to note that it leads to a disastrous theory of interpretation for Mohammed Ali. When Ali says "I am the greatest", what he's trying to get the rest of us to believe isn't the self-locating thing – he's not trying to get all of us to self-attribute *being the greatest*. What he wants to get across to all of us is *that Ali is the greatest*. (See Egan 2007 for the analogous problem with "my pants are on fire.")²⁵

²⁴ Officially, the view is that *all* (declarative) sentences have self-locating propositions (equivalently, properties) as their semantic values, and some (though not all) of them have as their contents self-locating propositions that don't just carve up the possible predicaments based on which world they're in. (Equivalently, properties that aren't just properties of inhabiting a certain sort of world.)

²⁵ Also, given the disastrous predictions of the self-locating theory of first-person indexicals, we already need a pragmatic story about the use of context-sensitive expressions to get people to believe self-locating things. For example, Ali's utterance above certainly *also* will succeed in getting credulous audience members to self-attribute the property, *being addressed by someone who is the greatest*. How does it do that? It does that because, when someone asserts a sentence that includes a context-sensitive expression, in order to accept the assertion, we need to take them to be in a position to express something true by uttering that sentence, and we need to take ourselves to be in a position to have something true expressed to us by an utterance of that sentence. And this will, of course, involve self-attributing some properties. But we needn't say that the

The second reason why I don't want to run a self-locating story here is that while (if we ignore the first worry) it might look initially attractive for the "you" cases, it looks really bad for pretty much all of the cases not involving "you". It sounds particularly awful for, e.g., the "ready" cases – it's not really clear what the relevant self-locating proposition would even *be*, exactly, and to the extent that I have a grip on how the proposal would go, it seems extremely unattractive. So if we want a general story about all of the cases involving context-sensitive expressions – and I think we should want this – we shouldn't go self-locating.

7. *Are You Talking to Me?*

Here is an apparent problem for the sort of line I've been pushing: Suppose Sven is listening in on Horton's sermon, and someone asks him for an evaluation of the truth or falsity of Horton's use of "Jesus loves you". Sven will base his evaluation on whether he thinks Jesus loves the members of the congregation Horton was addressing or not, not based on what he thinks about whether Jesus loves *him*. (The Robbins case is perhaps a better example here, since believing that the participants in the seminar can take control of their lives, though I can't, or that I can, though they can't, doesn't require any unorthodox religious views in the way that thinking that Jesus loves the members of the congregation, but not me (or v.v.) does.) In this respect, eavesdroppers' assessments of truth and falsity in cases like Horton's and Robbins's seem to pattern with paradigmatic uses of "you" in which it has a single clear referent (which might be an individual or a group) – such as when Daniels points at O'Leary and says "you are a fool." In the Daniels and O'Leary case,

properties that we need to self-attribute in order to take ourselves to be on the receiving end of a truth-expressing use of some sentence are the *semantic contents* of the sentences in question. So while it's certainly true that if Frank accepts Horton's assertion at the sermon, he'll come to self-attribute *being loved by Jesus*. But we needn't conclude from this that the property is the semantic content of Horton's utterance.

eavesdroppers will attribute truth and falsity based on whether they take O'Leary to be a fool, not on whether they take themselves to be fools.

On the face of it, this looks like a problem for the sort of shotgun-assertion theory I've been advocating, on which what's expressed to each audience member is the singular proposition about them. (And like a point in favor of the view that what's expressed is some proposition about a group.)²⁶ The fact that people who clearly aren't among Horton's intended addressees won't base their truth-value attributions on whether they take *themselves* to be among the objects of Jesus's love, is trouble for the view that what's expressed to everybody who hears or evaluates the utterance is the singular proposition about them. The thing to say, I think, is that there are some restrictions on who counts as an eligible addressee, and so on who gets to count as a genuine audience member for purposes of resolving context-sensitivity. (This is akin to some things that Stefano Predelli says in (Predelli 1996, 1998a, 1998b).) Since the eavesdroppers clearly weren't among the intended addressees, they don't get to count as audience members, and so they don't, strictly speaking, have anything expressed to them. Horton just wasn't talking to them, and so their context is not eligible to play the relevant semantic-value-determining role. Still, they're in a position to work out what Horton asserted to whom, and they'll base their truth-value ascriptions on whether they think those assertions were right or not. (Other kinds of cases that motivate constraints on eligible addressees, and in general on eligible audience contexts, are discussed in Predelli's papers cited above.)²⁷

²⁶ It's also another reason not to go for a relativist view here, since the eavesdropping data don't fall out in the way that motivates relativism about, for example, epistemic modals. (See e.g. Egan, Hawthorne and Weatherson 2005, Egan 2007 and MacFarlane forthcoming.)

²⁷ This points to a respect in which the sort of *audience* sensitivity that occurs here is different from the sort of *assessment* sensitivity MacFarlane discusses, since though anybody can be an assessor, not just anybody gets to count as an audience member. Some people just aren't being addressed. (This is non-disastrous since people

Conclusion

The central upshot of the above is that the speaker-positional theory of context-sensitivity is false: we should deny *SPEAKER ONLY* and accept *AUDIENCE SENSITIVITY*. Since there are cases where we want a speaker to express different contents to different audience members using a single sentence, we should say that the semantic values of context-sensitive expressions are sensitive not just to the properties of the speaker (at the time of the utterance), but also to properties of the various audience members. Once we adopt *AUDIENCE SENSITIVITY*, a couple of questions arise: Is this audience-dependence best accommodated by saying that there are sometimes more utterances being made than we might initially have thought, or by saying that the ratio of contents to utterances is higher than the 1:1 that we might initially have expected? And is it best accommodated by retaining a positional notion of context and saying there are two distinct contexts (that of the speaker and that of the audience member) playing a role in determining content, or by saying that contexts are something richer, and which context is fed into character isn't fixed just by the properties of the speaker? I've tried to say some preliminary things about those issues, but I take them to be much less central than the argument against the speaker-positional theory.

I'll close with a few remarks about the significance of accepting my diagnosis of the sorts of cases we've been discussing.²⁸ One interesting consequence is that we seem to lose first-person access to the contents of our utterances – if Horton doesn't know exactly who's in his audience, he won't be in a position to know just which propositions he's expressed by saying "Jesus loves you". Indeed, on this sort of view, Horton is liable to have asserted some

who aren't among the addressees can still do evaluation based on being able to figure out which contents they expressed to whom, even if strictly speaking, *no* content was expressed to them.)

²⁸ Much of what follows is due to extremely helpful discussions with Nicholas Silins and Yuri Cath.

propositions that he's not even in a position to *entertain*, since it's likely that there are some members of his audience to whom he lacks the sort of connection that's required to support singular thoughts about them. Because of this, we lose a certain tight connection between the contents of our utterances and our communicative intentions. And since the content that Horton's utterance expresses to me needn't be anything that he's in a position to entertain, we also can't quite retain the slogan I cited earlier, about content being "what travels" – as what moves, in successful communication, from the speaker's head to the hearer's. We should say instead that content is just what *arrives*.

All of this is, I think, interesting rather than catastrophic. We certainly don't need to sever the connection between content and speaker intentions altogether. Horton can, on the sort of view I've been arguing for, communicate to Frank the proposition *that Jesus loves Frank* without having, or being in a position to have, any intentions at all about that particular proposition. To do this, though, he does need to have a particular sort of general intention – an intention to communicate to each member of his audience the singular proposition about them.

References:

- Bach, Kent (1994) Conversational implicature, *Mind & Language* 9, pp. 124-162.
- Barwise, Jon and John Perry (1981) Situations and Attitudes, *The Journal of Philosophy* 78:11, pp. 668-691.
- Chisholm, Roderick (1982) *The First Person: An Essay on Reference and Intentionality*, Minneapolis, University of Minnesota Press.
- Egan, Andy, John Hawthorne and Brian Weatherson (2005) Epistemic Modals in Context, *Contextualism in Philosophy*, Preyer and Peter, eds., Oxford University Press, pp. 131-69.
- Egan, Andy (2006) Secondary Qualities and Self-Location, *Philosophy and Phenomenological Research* 72:1, pp. 97-119.
- Egan, Andy (2007) Epistemic Modals, Relativism, and Assertion, *Philosophical Studies* 133:1, pp 1-22.
- Egan, Andy (forthcoming) De Gustibus Non Disputandum Est (At Least, Not Always), forthcoming in *Disagreement*, Feldman and Warfield, eds., Oxford University Press.
- Kaplan, David (1989) Demonstratives, *Themes from Kaplan*, Joseph Almog, John Perry and Howard Wettstein, eds., Oxford University Press.
- Lasersohn, Peter (2005) Context Dependence, Disagreement, and Predicates of Personal Taste, *Linguistics and Philosophy* 28:6, pp. 643-686.
- MacFarlane, John (2003) Future Contingents and Relative Truth, *The Philosophical Quarterly* 53, pp. 321-36
- MacFarlane, John (2005) The Assessment Sensitivity of Knowledge Attributions, *Oxford Studies in Epistemology* 1, ed. T. Szabó-Gendler and J. Hawthorne, 197-233.
- MacFarlane, John (forthcoming) Epistemic Modals are Assessment-Sensitive, forthcoming in an Oxford University Press collection on epistemic modals, ed. B. Weatherson and A. Egan.
- Moltmann, Fredericke (forthcoming) Generic *One*, Arbitrary PRO, and the First Person, forthcoming in *Natural Language Semantics*.
- Lewis, David (1979a) Scorekeeping in a Language Game, *Journal of Philosophical Logic* 8, 339-59. Reprinted in David Lewis, *Philosophical Papers v.1*, Oxford University Press.
- Lewis, David (1979b) Attitudes *De Dicto* and *De Se*, *Philosophical Review* 88, 513-43. Reprinted with postscripts in David Lewis, *Philosophical Papers v.1*, Oxford University Press, pp. 133-160.
- Lewis, David (1980) Index, Context and Content, *Philosophy and Grammar*, Stig Kanger and Sven Öhman, eds., Dordrecht, Reidel. Reprinted in David Lewis, *Papers in Philosophical Logic*, Cambridge, Cambridge University Press (1998), pp. 21-44. Page references are to this reprinting.
- Liao, Shen-yi (MS) Time Travellers and Centered Worlds
- Perry, John (1979) The Problem of the Essential Indexical, *Nous* 13, pp. 3-21. Reprinted in John Perry, *The Problem of the Essential Indexical and Other Essays*, Stanford, CSLI Publications (2000), pp. 27-44.
- Predelli, Stefano (1996) Never put off until tomorrow what you can do today, *Analysis* 56:2, pp. 85-91.
- Predelli, Stefano (1998a) I am not here now, *Analysis* 58:2, pp. 107-115.
- Predelli, Stefano (1998b) Utterance, Interpretation and the Logic of Indexicals, *Mind & Language* 13:3, pp. 400-414.

- Sidelle, Alan (1991) The Answering Machine Paradox, *Canadian Journal of Philosophy* 21:4, pp. 525-539.
- Stalnaker, Robert (1978) Assertion, *Syntax and Semantics* 9, 315-332. Reprinted in Stalnaker 1999, pp. 78-95.
- Stalnaker, Robert (1998) On the Representation of Context, *Journal of Logic, Language and Information* 7. Reprinted in Stalnaker 1999, pp. 96-114.
- Stalnaker, Robert (1999) *Context and Content*, Clarendon, Oxford.