Assessment-contextual indexicals*

Josh Parsons

October 21, 2009

Abstract

In this paper, I consider whether tenses, temporal indexicals, and other indexicals are contextually dependent on the context of assessment (or *a-contextual*), rather than, as is usually thought, contextually dependent on the context of utterance (*u-contextual*). I begin by contrasting two possible linguistic norms, governing our use of context sensitive expressions, especially tenses and temporal indexicals (sections 2 and 3), and argue that one of these norms would make those expressions u-contextual, while the other would make them a-contextual (section 4). I then ask which of these two norms are followed by English speakers (section 5). Finally, I argue that the existence of a-contextuality does not in any sense entail "relativism" about truth (section 6).

1 Introduction

Traditionally, it has been supposed that each utterance of a natural language sentence expresses just one proposition — the same proposition to each person who hears it — and each proposition has just one truth value. Some recent work in the philosophy of language[1] casts doubt on these assumptions. It has been suggested both that assessment (and not just utterance) makes a contribution to context (so that one utterance may express different propositions to different hearers) and that truth may be assessment relative (so that one proposition may have different truth values for different hearers).

*Thanks to David Chalmers, Heather Dyke, Andy Egan, Michael Glanzberg, Jonathan Schaffer, Ian Spencer, and an anonymous referee for comments on and discussion of this paper.

[1]For example, MacFarlane (2003), Lasersohn (2005), Egan (2007).

Current supporters of assessment contextuality and assessment relativity present their arguments as descriptive semantics: they think that there are some patterns of usage in natural language (and particularly in English) that can be explained no other way. Defenders of the orthodoxy offer alternative explanations of those usages, and also, sometimes, philosophical arguments that assessment contextuality or (especially) assessment relativity are incoherent.[2]

It has sometimes seemed to me that the supposedly rival proposals offered by supporters of assessment contextuality and assessment relativity, and by supporters of the orthodoxy are dubiously distinct. For example, any assessment contextual theory on which a single utterance expresses different propositions to different hearers could be matched by an orthodox theory on which a single utterance expresses a single proposition, but where a single proposition is itself imagined to convey different information to different hearers. The assessment contextualist might retort that this approach misuses the term "proposition" — but could not the orthodox say that it's part of what she means by "proposition" that the one proposition per utterance rule be respected? What should be an empirical, descriptive, question about natural language has come to hang on a choice of theoretical framework.

In this paper, I want to try a different kind of defence of assessment contextuality. I aim to say what assessment contextuality is, in a way that makes it a straightforwardly empirical matter whether some given expression in natural language is assessment contextual (in sections 2 and 3). I link this characterisation of assessment contextuality with a more familiar and theoretical way of saying what it is (section 4). This should answer the charges that both that assessment contextuality is not distinct from utterance contextuality, and that it is incoherent. With that achievement in hand, I return to the descriptive question of whether any expressions of English are assessment contextual (section 5). Finally, I return to the question of assessment relativity, and argue that it cannot be given the same kind of defence I gave for assessment contextuality (section 6).

[2]For an example of an alternative explanation, see Glanzberg (2007). I'm not aware of anyone saying in print that MacFarlane, Lasersohn, or Egan are contradicting themselves. The assessment relativists of a previous generation were, however, confronted this way, or so I read Evans (1985) argument against "tense logic", and many of D.H. Mellor's arguments against "tense" (for which, see footnote 9). Some philosophers (Lewis 1980) also reject the distinctions whereby the relativists distinguish their theories from contextualism.

2 A creation myth for tense

Mammoth hunt case 1: Suppose that Joe and Kelly are hunting mammoths in the forest. Joe is hidden on the ground near a well-trodden mammoth path, clutching his spear. It is important that he hurls the spear just when a charging mammoth is about 20 meters away from him. Unfortunately, the forest is dense, and being on the ground, Joe will not see the mammoth in enough time to throw his spear, so Kelly, who has climbed a tree must warn him when she sees a mammoth approaching.[3]

Joe's and Kelly's problem is that they have not yet invented tensed language (and here I include temporal indexicals, and such words as "past", "present", and "future" among what is to count as tensed language). Their language has ways of saying "A mammoth is 20m away simultaneous with the sun reaching its zenith"; they even have a date system, so that they can say "A mammoth is 20m away at τ", where τ is a date. But none of these sentences will help Kelly communicate what she needs to when she sees a mammoth begin its charge. If she says, without reaching any prior arrangement with Joe, "A mammoth is 20m away simultaneous with the sun reaching its zenith", Joe will waste valuable time determining whether the sun is *now* at its zenith before throwing his spear, and will likely be trampled.[4]

There need be no suggestion that Joe's and Kelly's existing language is descriptively deficient. They both have the resources to give a complete description of the universe, let's say. What they are lacking is linguistic norms concerning *when* they are to make utterances. If indeed a mammoth charges at noon, and Kelly knows this, it would be as appropriate for her to say "A mammoth is 20m away simultaneous with the sun reaching its zenith" at the moment the mammoth charges as it would be ten minutes before or ten minutes after.

What Joe and Kelly need is a linguistic convention giving Kelly a sentence that it is only appropriate for her to utter *when* the mammoth is 20m away. When Kelly utters that sentence, Joe is to throw his spear. Moreover if Joe and Kelly had such a convention, no prior arrangement would be needed — knowledge of the convention gives Joe all the information he needs to throw his spear in a timely fashion when he hears Kelly's utterance.

[3]The stories of Joe and Kelly in this section are inspired by the various examples of the essentiality of indexicals and tenses given by Richard Gale (1962), Hector-Neri Castañeda (1989), John Perry (1979) and others. Joe is named after the unlucky machine-gunner of Gale's example.

[4]Joe and Kelly are not, of course, speaking English. They are speaking a prehistoric tenseless language, which I have here translated into English. I represent the tenseless verbs of their language with English present-tense verbs.

Suppose now that Joe and Kelly are the original inventors of tense. Faced with their problem, they agree to the following convention:

Convention u for "Mammoth!". (Instructions to utterers) Utter the sentence "Mammoth!" only at a time t such that a mammoth is 20m away at t. (Instructions to hearers) When you hear someone utter "Mammoth!", you may suppose that she is behaving in accordance with this convention.

Following this convention, Joe and Kelly are able to solve their problem. Kelly yells "Mammoth!" only when the mammoth is at the right distance, on pain of violating the convention, which she has no reason to do. Knowing this, Joe throws his spear when Kelly yells "Mammoth!" and kills the mammoth before it can trample him.

However, Joe and Kelly hunt game other than mammoth, and it would be useful to them to have a general convention that allows them to engage in this kind of communication in regard to other subject matters. They decide to adopt the following convention, instead of the special-purpose "Mammoth!" one:

Convention u. (Instructions to utterers) Utter a sentence Φ containing the word "now_u" only at time t such that $\Phi[t/now_u]$ — a sentence resulting from substituting all occurences of "now_u" in Φ with a date denoting t — is true.[5] (Instructions to hearers) When you hear someone assertorically utter a sentence containing "now_u", you may expect that she is behaving in accordance with this convention.

Following convention u, Joe and Kelly are able to solve their problems. Kelly can only say "A mammoth is 20m away at now_u" at the time at which the mammoth is 20m away, on pain of violating the convention, which she has no reason to do. Knowing this, Joe throws his spear when Kelly says "A mammoth is 20m away at now_u", and kills the mammoth before it can trample him. Tense is born.

3 Rival conventions

Convention u is not the only convention Joe and Kelly could have adopted. They could instead have chosen this one:

[5] It might be objected that Joe and Kelly would "waste valuable time" computing the substitutions while following this convention just as they would if they spoke a tenseless language. This, however, misses the point of the earlier "waste valuable time" objection. There the problem was not that time would be lost producing or understanding utterances, but that the conventions of a tenseless language give no guidance at all as to when any sentence should be uttered.

Convention a. (Instructions to hearers) For any time t at which you hear someone assertorically utter a sentence Φ containing the word "now_a", adjust your beliefs as you would had you heard the sentence $\Phi[t/now_a]$. (Instructions to utterers) When someone hears you utter a sentence containing "now_a", you may expect that he will behave in accordance with this convention.

I wish to make three important points in regard to convention *a*. First, there is nothing incoherent about convention *a*. It represents a linguistic convention that could govern tense (or, to be more careful, could govern a tense-like feature of language). Second, convention *a* is in no way equivalent to convention *u* (though of course, they are similar). And third, we ought to expect to see something like convention *a* in natural languages.

3.1 Convention a is coherent

It's hard to say more about the coherence of convention *a* without saying more about why someone would think that it is incoherent. And it's hard to do that without connecting it with another view about tense that I will introduce later (in section 6. To foreshadow slightly, convention *a* has something in common with views that Arthur Prior (and others) held about tense, and which have excited vigorous opposition on both metaphysical and semantic grounds.[6]

My principal argument that convention *a* is coherent is that, described as I described it, it's obvious that it's coherent. There are other ways of describing it to which people object that they are incoherent; however, when it is seen that these other descriptions are equivalent to convention *a*, the objections fall away. More on this below.

3.2 Convention a is distinct from convention u

The difference between convention *a* and convention *u* comes out when we imagine cases where some time elapses between when a sentence is uttered, and when it is heard.[7]

[6]From, for example, D.H. Mellor and Gareth Evans. For more on Mellor's objections, see footnote 9 of this paper. For Evans's, see his (1985), and my (2003).

[7]This point was apparently first made by Keith Donnellan, as reported by David Kaplan (1989, p. 491, footnote 12). Donnellan went on to suggest, as I do, below, that there could be a language with two distinct temporal indexicals, one governed by convention *u*, the other by convention *a*. Unfortunately, in his short discussion of this point, Kaplan doesn't clearly distinguish Donnellan's imagined use of "now" (in which "now" refers to the "time of audition") from the very different idea that "now" recorded on an answering machine refers to the "time of playback". Donnellan's

Mammoth hunt case 2: Suppose that Kelly is quite a long way from Joe — too far for him to hear her directly — but that the walls of the valley they are in are such as to create an echo of anything Kelly says, which Joe will hear 5 seconds after Kelly says it. Suppose Kelly sees a mammoth at 25m distance. It would then be appropriate for her to yell "A mammoth is 25m away at now_u". But it would *not* be appropriate for her to yell "A mammoth is 25m away at now_a", because by the time Joe hears her, the mammoth has charged up to 20m distance. Convention *a* says that Kelly should expect that Joe will come to believe that the mammoth is 25m away *when he hears her utterance*, and she knows that when he hears it, the mammoth will not be at 25m distance. Kelly would be willfully misinforming Joe — lying to him — if she said "A mammoth is 25m away at now_a".

This shows that "now_u" and "now_a" are not synonymous; and that, since they were stipulatively introduced by conventions *u* and *a* respectively, the two stipulations did not have the same effect; so the conventions are not equivalent.

In the previous example, Kelly could use either "now_u" or "now_a" to do what she wants, namely, to help Joe. She knows that the mammoth will charge 5 meters in 5 seconds, and that it will take 5 seconds for her utterance to reach Joe. She could say "A mammoth is 25m away at now_u" when she sees that the mammoth is 25m away, or she could say "A mammoth is 20m away at now_a" when she sees that the mammoth is 25m away, judging that the mammoth will reach the 20m mark just when Joe hears her.

Mammoth hunt case 3: Suppose now that Kelly knows how fast the mammoth is charging, but that Joe does not. This would make "now_a" much more useful to Kelly than "now_u". If, when Kelly sees a mammoth at 25m distance, she yells "A mammoth is 25m away at now_u", that leaves Joe to figure out where the mammoth is now, based on his own judgement of when Kelly yelled. In the scenario as we are now describing it, Joe is not well placed to do that, because he doesn't know how much ground the mammoth will cover in the time it takes Kelly's utterance to get to him. But, in contrast, if Kelly yells "A mammoth is 20m away at now_a", Joe knows that she has already taken into account the time it takes for the mammoth to get to him. In this scenario, it would be very useful to both Kelly and Joe to have "now_a" in their language in addition to "now_u".

---

imagined "now", as I understand is, is a form of assessment contextuality; Kaplan's take on the answering machine use of "now" is not. As it happens, I agree with Kaplan about the answering machine case; "now" recorded on an answering machine is *not* assessment contextual — as can be seen in a case where there is a delay between playback and assessment (see section 5).

3.3 We ought to expect to find something like convention a in natural languages

I think that there is an in principle gap in a language that lacks "now_a" (and lacks anything like it). In principle, Kelly could meet situations that can be dealt with better by using "now_a", than by using "now_u". It's not clear that the case described above is such a situation. Since the problem is that Joe doesn't know how fast the mammoth is charging, couldn't Kelly just yell "A mammoth is 25m away at now_u and is charging at 1 meter per second"?

Mammoth hunt case 4: Suppose now that Kelly knows how fast the mammoth is charging, that Joe does not, and moreover, that Joe is a bit slow at the best of times, poor at simple arithmetic, and that Kelly has just seen a falling branch knock Joe on the head. Joe appears dazed and might not be able to draw the inferences needed to take timely action when he hears "A mammoth is 25m away at now_u and is charging at 1 meter per second", even if he knows all the relevant information. But if he hears "A mammoth is 20m away at now_a", then he needs to perform fewer inferences in order to take timely action. Kelly cannot achieve the same effect by using "now_u" and, in addition, giving Joe extra information about the mammoth's speed. That would only add to the number of inferences Joe must perform correctly, and only make it harder for Kelly's attempt at communication to succeed.

The case just described is a situation in which it would help Kelly to be able to use "now_a" rather than "now_u". But the relationship between "now_u" and "now_a" is symmetrical — there are scenarios in which it would greatly help Kelly and Joe to have "now_u" in their language, and in which "now_a" would be little help.

Mammoth hunt case 5: Suppose now that Joe knows how fast the mammoth is moving, and is arithmetically competent. A branch has just fallen on Kelly's head, she is dazed and can't really tell how fast the mammoth is moving, all she can remember is that she was supposed to warn someone about a charging mammoth. She sees the charging mammoth at 25m. She has neither the information nor the inferential competence to responsibly say "A mammoth is 20m away at now_a", but she can yell out "A mammoth is 25m away at now_u", and leave whoever hears her to draw the appropriate inferences.

Finally, we might suppose that Joe and Kelly set out on their mammoth hunt without knowing what will befall them; who will be best placed to judge the mammoth's speed; on whose's head a branch might fall. It would be to their benefit to have both conventions. This leads to an interesting question: do we English

speakers have both conventions? (If not, do other natural language speakers?) We ought to expect that they do have *something like* both — or at least, it would be surprising and in need of explanation if natural languages had *nothing* remotely like convention *a*. This is a question I will return to in section 5.

4 Character semantics

I now want to consider a very different way of distinguishing "now_u" from "now_a" — a way that has received some attention in the literature. "now_u" is a very standard Kaplanian indexical. Different tokens of it refer to different things, depending on when they are uttered. But there is a stable meaning to the word, which all the tokens have in common — all tokens of "now_u" depend on their context of utterance in the same way — all refer to the time when they were uttered.

Let us remind ourselves of Kaplan's (1989) terminology: expression types have *characters*, which together with a *context* of utterance determine a *content*. In the case of sentences, contents are propositions; in the case of singular terms, contents are their referents. Characters are rules for getting the content of an expression as used — functions from contexts to contents. The character of an expression is its meaning — it is what two semantically equivalent expressions have in common — the content is "what is said" by the expression — what a listener should take from it.

The character of "now_u" is the function that maps contexts of utterance of the word "now_u" to the time of that context. The character of the sentence "A mammoth is 25m away at now_u" is the function that maps contexts of utterance to propositions that say that a mammoth is 25m away at the time of utterance that is part of each context.

This semantic account of "now_u" entails convention *u* via a general norm of truthfulness:

Convention of truthfulness. (Instructions to utterers) Utter a sentence assertorically only at such a time and in such a way that your utterance comes out true. (Instructions to hearers) When you hear someone assertorically utter a sentence, you may expect that she is behaving in accordance with this convention.

It follows from the specification of the the character of "now_u" given above that Kelly's utterances of sentences containing "now_u" will only be true if the rest of the sentence is satisfied by the time at which she does the uttering. So if Kelly follows the convention of truthfulness, then she will utter a sentence Φ containing

"now_u" at such a time and in such a way that a sentence like Φ but in which "now_u" has been replaced by a name for the time at which she does the uttering comes out true. That is to say, if Kelly follows the convention of truthfulness then she follows convention *u*. Given the character of "now_u", convention *u* is a special case of the convention of truthfulness.

But "now_a" cannot be given a standard Kaplanian semantics, because convention *a* cannot be derived from the convention of truthfulness plus a specification of character. (Not at least, if a character is, as standardly thought, a function from contexts of utterance to contents). This is because what convention *a* asks hearers to do depends on *when they hear* an utterance of "now_a". Convention *a* could ask two hearers to do different things on hearing one and the same utterance. But what the convention of truthfulness asks utterers and hearers to do depends only on the truth conditions — the content — of the utterance in question. On the Kaplanian semantics, each utterance has once-and-for-all truth-conditions, which do not depend on when the utterance is heard.[8] So the instructions to hearers that you might derive from convention of truthfulness plus a Kaplanian character cannot depend on when the utterance is heard.

A simple revision to Kaplan, however, will accomodate "now_a". What we should say about sentences containing "now_a", I think, is that they say different things to different people. The content of — what is said by — such a sentence depends not only on its context of utterance, but its context of assessment. So we should expand the Kaplanian notion of *context* to include features of the context in which a sentence is assessed, not just the context in which it is uttered. Along with this goes an expansion of the notion of *character* of course. A character now becomes a function from pairs of contexts of utterance and assessment to contents. The Kaplanian notion of *content* remains unchanged, still corresponding to the traditional proposition in the case of sentences, and to the referent of a singular term.

Given this specification of the character of "now_a", the convention of truthfulness will entail convention *a*, because if Kelly follows the convention of truthfulness, Joe can expect that she will only utter "A mammoth is 20m away at now_a" at such times and in such a way that his assessment of her sentence comes out true. That is, when he hears her, he can trust that a mammoth is 20m away at the time of that assessment. So the convention of truthfulness asks Joe, the hearer, to behave as convention *a* asks him to. Since Kelly can then expect that Joe will behave as the convention of truthfulness asked him to, she can then expect that he

[8] A slight complication: Kaplan's own discussion of character semantics makes the content of a sentence a function from a world / time / individual triple to a truth value. Some people have interpreted this as a form of relativism (see section 6) — as a denial that propositions have once-and-for-all *truth values*. That's not what I'm presenting Kaplan as denying here.

will behave as convention *a* asks him to, which is in turn what convention *a* asks her to do.

To introduce some terminology, an expression is *contextual* iff it can have more than one different content in different contexts — that is, iff its character's codomain has more than one member. Following MacFarlane (2003), we might say that an expression is *u-contextual* iff it is contextual and its content depends only on its context of utterance, and *a-contextual* iff it is contextual and its content depends only on its context of assessment. Finally, a sentence is of *mixed contextuality* iff it is contextual, but neither *u-contextual* nor *a-contextual*. Described in these terms, we have just been noticing that "now_u" is u-contextual — it is a standard indexical — and that "now_a" is a-contextual. A sentence containing both would be of mixed contextuality.

MacFarlane and others have been anxious to defend a-contextuality. I am joining their ranks, though not necessarily agreeing with their applications of a-contextuality. (Nor am I agreeing with their more radical view that there is, in addition to a-contextuality, an even weirder phenomenon — "assessment relativity". More on that later). In my view, a-contextuality is an important tool to have in our semantic kit, and we should be on the lookout for it in natural language.

I also have an argument, distinct from those of MacFarlane, that a-contextuality is coherent, and that it is distinct from u-contextuality. You've now heard pretty much all of that argument. I showed in section (3) that convention *a* is coherent. There's nothing to stop two people agreeing to use words that way; nothing to stop a whole linguistic community doing so. I also showed that convention *a* is distinct from convention *u*, and, in this section, that convention *a* corresponds to a-contextuality in our modified Kaplanian framework. Since there's nothing incoherent about convention *a*, there's nothing incoherent about a-contextuality.

5 A-contextuality in English

So far I have been putting off the task of saying whether there really is any a-contextuality in English. In particular, I put off the task of asking whether the English word "now" is more like "now_u" or like "now_a". I will argue that "now" is more like "now_u", but that there is *some* evidence of words that work like "now_a" in English. I am less certain of the conclusions that I reach in this section than of the conclusion of the previous sections. It seems clear to me, as a matter of philosophical theory building, that there's nothing incoherent about a-contextuality. In contrast, matters of the linguistic norms in force in English are more empirical, and I am not going to attempt any empirical linguistics in this paper. As philoso-

phers often do when studying language, I present some explanations of my own linguistic intuitions in the hope that the reader will share them.

5.1 "Now" and the answering machine paradox

Returning to the question of which of Joe and Kelly's two indexicals most resembles "now", there is one respect in which "now" is unlike both. In Joe and Kelly's language "A mammoth is 20m away at now_u" is perfectly grammatical, but "A mammoth is 20m away at now" is not grammatical in English. This is a philosophically uninteresting difference between "now" and my made up words. It's just that "now" doesn't have an accusative case, so it can't be used in contexts that demand one in English.

The question I want to ask is not, is "now" exactly like "now_u" or exactly like "now_a", but does "now" refer to the time of the context of utterance, or the time of the context of assessment? It seems clear to me that the primary use of "now" is to do the former. To see why, imagine a journalist making a recorded report from a war zone: "Now the invasion is just beginning." Later the report is broadcast, by which time, of course, the invasion has begun some time ago. If "now" were a-contextual, this would be poor journalism — our journalist would be misreporting the facts. On the contrary, this is a perfectly acceptable and standard use of the word "now". There's not even anything jokey about the journalist's use. This shows that "now", as normally used, refers to the time of utterance, not of assessment.[9]

The so-called "answering machine paradox" is the famous (alleged) counterexample to this. I record a message on my answering machine that says "I'm not here right now, please leave a message". It seems as though, insofar as what I say is true, the word "now" has to be construed as a-contextual, as referring, in your context of assessment, to the time at which you hear the message. This is an interesting case because there is something standard and conventional about it — people do leave messages like that. But even non-philosophers, I think, hear a joke in doing so. The joke is a pun on two meanings of the word "now", the primary u-contextual meaning, and a secondary a-contextual meaning. So, you

[9]This type of argument perhaps originates with D.H. Mellor, who has used something like it in a number of places. In (1998, pp. 78–79), for example, he says "No one thinks... that my death will posthumously verify every premature announcement of it". Mellor is perhaps over-confident of the scope of this point — he seems to regard it as apriori, no matter what language the announcements are made in. In my view, whether premature announcements of Mellor's death will be true as read by attendees at his funeral depends on the linguistic norms governing the language in which those annoucements were made. But I think that Mellor was onto something important about tense in English here, on the assumption that the annoucements were made in English.

might think, though "now" is usually u-contextual, it can have an an a-contextual meaning, if the conversational circumstances are right.

Unfortunately, however, a-contextuality alone can't explain what's odd about the answering machine paradox. Suppose that I ring you, and your answering machine answers. However, there is a delay on the line, and in the intervening time between answering machine playing and my hearing the message, you have arrived home (and let us suppose, if it makes a difference, that I know that there is such a delay). Is the proposition expressed to me by the answering machine message true or false? Have you lied to me by not taking enough care that I would not hear the message at an inappropriate time? I find it hard to have an intuition on this, but many philosophers I've spoken to feel strongly that the message is true under these circumstances. But then the time that "now" indicates is not the time of reception by the message's hearer but the earlier time at which the answering machine played the message. [10]

5.2 "You" and the road accident case

Where could we look for an example of a word whose primary meaning is a-contextual? Andy Egan has convinced me[11] that "you" is a good candidate, though the argument I give below is my own:

Road accident case: Suppose I am standing on a crowded sidewalk, and I see someone I don't know — call him Mr. X — step into the road. X is looking in the wrong direction, and will be hit by a car unless he steps back immediately. I need to warn him, and I'm not close enough to just grab him and pull him back from the road. Obviously enough, I can warn X by communicating with him. So I might yell "Hey you, look out! You're about to be run over!"

[10] Thanks to Jonathan Schaffer for urging this point on me. The intuition that "now" in the answering machine paradox refers to the time of playback, rather than of assessment, seems to have been shared by no less an authority than David Kaplan (1989, p. 491, footnote 12). (On this passage in Kaplan, see also my previous footnote 7). If this intuition is part of the meaning of "now", as used in the answering machine case, then it could suggest a u-contextual explanation of the case along the "pretence" lines suggested below. Perhaps "now" is always u-contextual, but in the answering machine scenario, the hearer pretends that the speaker is speaking the words at the time they are played back. Another u-contextual explanation of the answering machine case is Andy Egan's idea of an "utterance bomb" — that we should think of the recorded message as an utterance that takes place at a spatial and temporal distance from the utterer. The delay between playback and audition case described here in the text seems to favour either u-contextual explanation of the answering machine paradox over an a-contextual one.

[11] In a paper given at the Australasian Association of Philosophy conference 2006, and in personal communication.

It seems to me that what happens when I do this is that everyone who hears me, and follows the linguistic conventions of English, straightaway starts to form the self-locating belief that they are about to be run over, and takes appropriate action. Many of them quickly realize that the belief they're forming is false, of course. But my utterance gets the message across to Mr. X immediately — he doesn't have to make any inferences about my intentions in order to form the self-locating belief. And that's just what you'd expect if "you" was a-contextual, and referred to the person assessing the utterance of "you". In effect, I've used "you" to express a different Kaplanian content to each person who hears me. To Mr. X, I've said that Mr. X is about to be run over, and to Ms. Y, standing safely on the sidewalk, I've said that Ms. Y is about to be run over. I've said something false to most of the people who've heard me, but, after all, saving a life is more important than following the norms of assertion perfectly.

Consider, by way of contrast, what would have to happen if "you" was u-contextual. If it were, then it would have to be that when I said "You're about to be run over!", everyone who hears me gets the same proposition by interpreting my utterance. Since Mr. X is able to infer from what I've said that he's about to be run over, it must be that I've said something general; someone which I can say to everyone, from which each person can infer their own case. That is, "you" is not functioning as a referring expression like "I", but as a variable, like "someone", ranging over a u-contextually specified set of people that includes Mr. X. There are lots of different ways this could work. Perhaps the content of my utterance is more like "Everyone I intend to communicate with is about to be run over"; perhaps it is more like "Someone who hears this is about to be run over".

What they all have in common is that they are poor ways to communicate with Mr. X. I want to get Mr. X to form the self-locating belief that he is about to be run over; and I want to get him to form it as quickly and efficiently as possible, given that he is a complete stranger to me. He cannot do this if he has to make inferences from my communicative intentions, which are unknown to him. Nor can he do it if the content of my utterance is something that is clearly false such as "Everyone who hears this is about to be run over." But, I think, not only would I like to communicate the self-locating belief to Mr. X by using the word "you", should the situation arise, using it would be effective. I can't imagine how any content that is sufficiently general to be expressed to *every* hearer could do the trick.

Notice by the way, that the problem that "Everyone who hears this is about to be run over." is false is quite different from the situation we have with my a-contextual treatment of "you", where I say something false to everyone who hears me except Mr. X. What would make it impossible to communicate effectively

with Mr. X is if he has to infer something from the content of my utterance, when he has every reason to suppose that that content is false. Saving lives is more important than speaking truly, but if I am to save Mr. X by communicating with him, I must say something that he can believe.

In discussing the mammoth hunt cases, I said that the absence of "now_a" from Joe and Kelly's language would be an "in principle gap". There are situations which Joe and Kelly could face in which "now_a" would help them, but in which "now_u" could not. If I am right in my treatment of the road accident case, then in English this gap seems to filled by the second person.

5.3 "You" and the recruiting poster case

Another type of argument that "you" must be a-contextual appeals to utterances that are broadcast to many people. This is the type of argument Egan used:[12]

Recruiting poster case: Suppose that the Syldavian army have produced a propaganda poster urging citizens to join the armed struggle against Borduria. The poster features a picture of a familiar national military figure pointing as if out of the picture towards the reader, and the words "Syldavia needs *you*!"

I've never been the target of such a recruiting campaign, but I'm familiar enough with how the sentence on the poster is intended to affect its readers, and why it would be effective. Seeing the poster, the reader's mind is immediately drawn to the self-locating belief that Syldavia needs them. And this is just what we should expect if "you" is a-contextual. The content of the inscription on the poster depends on who's looking at it — to each person, it says that Syldavia needs that very person.

Again, consider what would have to happen if "you" is u-contextual. As before, in the road accident case, it must be that the content of "Syldavia needs you" is the same to each reader, and so it must be sufficiently general that each reader can infer that he or she is needed by Syldavia. The poster's message could be effective if this was what it meant, but no more effective than a poster which says "Syldavia needs *everyone*" or "Syldavia needs *all Syldavians*", or even "Syldavia needs *everyone who reads this*". Whereas in fact it seems obvious that the phenomenology of viewing a poster that says "Syldavia needs you!" is quite different from that of viewing a poster that quantifies over people whom Syldavia needs —

[12]But my own example and formulation. Any obvious flaws in the argument, unless otherwise noted, are mine, not Egan's.

when you see a poster like that uses the word "you", you *immediately* think of Syldavia needing *you*; you don't infer this from any generality about whom Syldavia needs. If you did do that inferring your mind would be drawn to thoughts that the poster designer hoped it wouldn't be drawn to, about how small your individual contribution to the war effort will be.

There is, however, another account of what's going on in the poster case that mimics the results of the view that "you" is a-contextual. It might be thought that when I understand a piece of written language, we do so by simulating the experience (or the cognitive effects of the experience) of having someone say the words to me. So the idea is that when I read "Syldavia needs you!" it's like I pretend that someone is saying "Syldavia needs you!" to me. I'm not endorsing this view, it's just that I see how a philosopher might come up with something like that. Call it the *simulation theory of writing*.

Whatever the other merits or demerits of the simulation theory, it fits well with the phenomena of the recruiting poster, even on the assumption that "you" is u-contextual. The simulation theorist says: what gives the poster its effect is not just the words, but the picture of the national military figure pointing as if out at the viewer. That image is a *prop* in Walton's (1990, pp. 37–38) sense. It invites me, as I look at the poster to not just pretend that someone is saying "Syldavia needs you!" but to pretend that someone is saying that while pointing and gazing intently at me. The words just mean "Syldavia needs everyone who is intended to hear this", but the gesture leaves no doubt in my mind as to who exactly is intended.

For this reason, I don't think that recruiting poster cases are the best argument for the a-contextuality of "you". It's not that I think that the simulation analysis is clearly better than an a-contextual approach — I just worry about relying too closely on examples that make essential use of written language. By contrast, my road accident case involves only spoken language, and so is a safer guide to the semantic properties of that language.

6 Relativism

Utterences that contain an occurence of some a-contextual expression can express different contents to different people who hear them. So, it can happen that one and the same token utterance is true for one hearer and false for another. If I am right about "you", for example, then, in the road accident case, when I yell "You're about to be run over!", that very utterance is true for Mr. X, and false for other people who hear it. Now that I have introduced you to "now_a", I can leave

an answering machine message saying "I am not at home now_a" that is false when I hear it, and becomes true when you hear that very same token sentence.

This might seem like a kind of objectionable relativism. What's true for me isn't true for you. But there's nothing more objectionable to this than there is to the familiar idea that different utterances of one and the same u-contextual sentence might have different truth values. Both are safe and unobjectionable because both are the same phenomenon of the content of a sentence varying from context to context.

Some philosophers hanker after a more robust type of variation of truth value. One traditional idea is that, when the world changes, propositions themselves should change their truth values — the doctrine of *proposition temporalism* or of *Stoic propositions*. (Bigelow 1996) That's different from a-contextuality. An a-contextual sentence expresses different propositions to different hearers, but says nothing about the nature of those propositions. They could be Stoic propositions, something like sets of world / time pairs, or they could be old-fashioned eternalist propositions, something like sets of worlds; or they could be Russellian propositions, or they could be some variant on Russellian propositions that allows for one proposition to have different truth values at different times.

Another idea that has gained ground recently, particularly in the work of John MacFarlane (2003), Peter Lasersohn (2005), and Andy Egan (2007), is that propositions should have different truth values in relation to different individuals, in the same way that Stoic propositions are supposed to have different truth values in relation to different times.[13] This does sound like an objectionable form of rel-

[13]As I say in the text below this footnote, this paper is neutral on whether the arguments for and against relativism cited above succeed. It seems appropriate, however, to briefly survey the types of evidence given in favour of relativism about truth, in order to emphasise the contrast I am drawing between a-contextuality and the type of truth-value variation that relativists believe in.

There seem to be three types of purported evidence for relativism. First, there are arguments from metaphysical premises about the nature of time and change; these are traditional arguments for temporalism about propositions. Second, there is linguistic data about what English speakers will say about truth (e.g. "What he said is true!") and thus about whether and how truthbearers can vary in truth values. Thirdly, there is the idea that discourse concerning what non-philosophers might loosely call "subjective" matters demands an analysis in terms of relative truth. Roughly, these are the arguments advanced by MacFarlane, Egan, and Lasersohn respectively.

Arguments one and three have a bit more of a history. Argument one seems to be at the root of much traditional critique of the "tenseless theory" of time. Gale's (1962) seems to me to be arguing in favour of a kind of relativism that tries to eliminate the "burden of inference" I describe in the text, below. (It is in honour of this paper that I have named one of my characters "Joe"). Something like argument three must have been put forward by first year students of philosophy since Plato's time, though not so capably as by Lasersohn.

I don't find any of these arguments decisive. It seems to me, in particular, that sufficiently sophisticated contextual approaches can tackle the phenomena described by two and three. In the

ativism to many people, and its proponents are hard at work to convince others that it is not. This too, of course, is simply a different and independent idea from a-contextuality (a point that MacFarlane has made clearly).

For the purposes of this paper, I have no brief either way on Stoic propositions or on MacFarlane-style relativism. However, I am more convinced that a-contextuality is in no way objectionable than I am that relativism is. The argument of section 3 convinces me that a-contextuality is a coherent way that a bit of language could work, and that it is distinct from u-contextuality. There's nothing impossible about a linguistic community agreeing to adopt convention *a*, and a linguistic community that adopts it is speaking a language that has some a-contextual expressions. In contrast, it is not obvious to me that there is some, coherent, independently describable, linguistic norm that we could adopt that would make our language truth-relativist.

There is a sort of norm that does seem close to what the relativists want. To motivate it, think back to the mammoth hunt cases. When I hear a u-contextual sentence, I have to do some work before I form the appropriate belief — I have to figure out what relation the context of utterance bears to me. For example, when Joe hears Kelly say "A mammoth is 25m away at now_u", he has to figure out what time Kelly said that (or risk making a mistake). A-contextuality just shifts this burden of inference to the speaker. Before Kelly says "A mammoth is 20m away at now_a", she has to think about who is going to hear her, and when, and adjust what she says appropriately (or risk miscommunicating). I call this phenomenon the *burden of inference*. (Of course, I'm not claiming that we actually perform these inferences — usually, we risk miscommunicating).

It sometimes seems that what relativists are after is a way of speaking that removes the burden of inference entirely — by allowing me, as it were, to pick up a self-locating belief from my head and drop it into yours. Suppose we took the "Instructions to utterers" part of convention *u*, and the "Instructions to hearers" part of convention *a*, and pasted them together, to make something like this:

Convention r. (Instructions to utterers) Utter a sentence Φ containing the word "now_r" only at time t such that $\Phi[t/now_r]$ — the sentence resulting from substituting all occurences of "now_r" in Φ with a date denoting t — is true. (Instructions to hearers) For any time t at which you hear someone assertorically utter a sentence Φ containing the word "now_r", believe true the sentence $\Phi[t/now_r]$.

A community that follows this convention would not suffer from the burden of inference — at least, not in regard to sentences containing "now_r". Kelly can

case of one, I just find the metaphysical premises completely unappealing.

say "A mammoth is 25m away at now_r" without having to think about who will hear her, and Joe doesn't have to think who's speaking, and when, to understand her. Also, convention r is coherent, at least in the sense that it does not entail a contradiction. But of course, grave miscommunication would take place in a community that accepted convention r. Convention r instructs Kelly to say "A mammoth is 25m away at now_r" when the mammoth is 25m away, and Joe to believe that the mammoth is 25m away whenever he hears her say that, even if some time has elapsed. In the case in question, that would lead to Joe's being trampled. Convention r instructs the members of a linguistic community to systematically miscommunicate.

Egan calls this misfeature — which is also a consequence of his version of relativism — the problem of "disasterous assertions". (2007, pp. 11–13) He argues that the particular applications for which he intends a relativist semantics will not suffer from it. He may be right about that. My point is that my argument that a-contextuality is a perfectly coherent possible linguistic phenomenon does not carry over to show the same about relativism. If, in fact, we do not speak a language that contains any a-contextual expressions, that seems to be a matter of historical contingency — we could have used convention a, but in fact we do not. If, in contrast, we do not speak a language governed by anything like convention r, it is hardly surprising, given that convention r directs us to make disasterous assertions.

Notice it does not help matters to suppose that users of "now_r" take care only to use this expression under circumstances under which disasterous assertions would not occur — where there the gap between the time of an utterance of "now_r" and the time of its reception is insignificant for example. For that "taking care" would reintroduce the burden of inference, and remove the distinctive relativistic features of now_r. There are two ways that this could happen, depending on whose responsibility it is to avoid disasterous assertions. Suppose that speakers take care never to utter a sentence containing "now_r" in a way that would lead to miscommunication; then the burden of inference would rest with the speaker, and now_r would be equivalent to "now_a". Suppose instead that speakers utter sentences containing "now_r" in any way they like provided it is consistent with convention r, but that hearers take care not to pay attention when miscommunication would result; then the burden of inference would rest with the hearer, and now_r would be equivalent to "now_u".

7 Conclusion

Recent discussions of context have turned up a variety of non-standard ways of thinking about natural language expressions that are normally regarded as contextual. I have argued that one such way at least, which I call "a-contextuality", is coherent, and is distinct from the way we ordinarily think about contextual expressions. It also seems to be a feature of language that could, in principle, be useful in communicating with other people when that communication is affected by a time delay.

It is another question whether any natural language is a-contextual. I've presented some evidence that there is a-contextuality in English. There may be other ways of explaining that evidence — but even if this is the case, it would be a striking fact, in need of explanation in itself, if there were no a-contextuality in any natural language.

References

Bigelow, J. (1996). Presentism and properties. *Philosophical Perspectives 10*, 35–52.

Egan, A. (2007). Epistemic modals, relativism, and assertion. *Philosophical Studies 133*, 1–22.

Evans, G. (1985). Does tense logic rest upon a mistake? In *Collected Papers*. Oxford: Clarendon.

Gale, R. (1962). Tensed statements. *Philosophical Quarterly 12*(46), 53–59.

Glanzberg, M. (2007). Context, content, and relativism. *Philosophical Studies 136*, 1–29.

Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, and H. Wettstein (Eds.), *Themes from Kaplan*. Oxford: Oxford University Press.

Lasersohn, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy 28*, 643–686.

Lewis, D. (1980). Index, context, and content. In *Papers in Philosophical Logic*, pp. 21–44. Cambridge: Cambridge University Press.

MacFarlane, J. (2003). Future contingents and relative truth. *Philosophical Quarterly 53*, 321–336.

Mellor, D. H. (1998). *Real Time II*. London: Routledge.

neda, H.-N. C. (1989). *Thinking, Language, and Experience*. Minneapolis: University of Minnesota Press.

Parsons, J. (2003). A-theory for tense logicians. *Analysis 63*(1), 4–6.

Perry, J. (1979). The problem of the essential indexical. *Noûs 13*(1), 3–21.

Walton, K. L. (1990). *Mimesis as make-believe*. Cambridge, MA: Harvard University Press.