

The Conversational Role of Centered Contents

Max Kölbel

For *Inquiry*

Abstract: Some philosophers, for example David Lewis, have argued for the need to introduce *de se* contents or centered contents, i.e. contents of thought and speech the correctness of believing which depends not only on the possible world one inhabits, but also on the location one occupies. Independently, philosophers like Robert Stalnaker (and also David Lewis) have developed the conversational score model of linguistic communication. This conversational model usually relies on a more standard conception of content according to which the correctness of believing a content depends merely on the possible world one occupies. The aim of this paper is to develop a modified conversational score model that operates with centered contents. I begin by explaining how in principle centered contents can figure in the transfer of information from one thinker to another. Here I distinguish the local portability approach from the portable surrogate approach. Then I explain how these modes of information transfer can be exploited in a stalnakerian conversational model involving centered contents, proposing a modified update rule for assertion.

1. Preliminaries¹

Semantic theories for natural languages usually employ some notion of semantic content. A semantics for a given language *L* will assign such contents to sentences at contexts of use. These assignments are supposed, via some further assumptions, to allow predictions as to the conditions under which the utterance of a sentence of *L* would be true, and derivatively predictions about logical truth and logical consequence (see e.g. Lewis 1970, p. 190; Kaplan 1977, p. 522; Predelli 2006; Stanley 2007, introduction; Partee 2010). Data concerning the truth conditions of (potential) utterances are assumed to be observable² independently and thereby to provide a method for empirically testing the semantics. The contents assigned to sentences in context are usually treated as multi-purpose representational contents, i.e. contents that can model representational aspects of linguistic utterances as well as of mental states, such as beliefs. Whether they are called “thoughts”, “propositions”, “contents” or whatever, the basic idea is the same: these entities represent the world in some way, and certain utterances (e.g. assertions) or states (e.g. beliefs) with such an entity as content can be assessed as to their success in representing the world. This

¹ Those already familiar with recent discussions of communication with centered contents, who are keen to cut straight to the issues, and who are not interested in preliminary or foundational questions, can safely skip this section.

² Observation is a problematic notion, and it is no less problematic in this case. If, for example, we think of the predictions as concerning a shared language, then we might consult the judgements of competent users of that language as to the conditions under which it would be correct to use a sentence. Thus we would be operating on the assumption that there is a correlation between these users’ judgements regarding the correctness conditions and the correctness conditions. This assumption may be a kind of stipulation (when we say what we mean by a “shared language”, we might say things that entail the correlation). But the problems don’t end here. We still need to separate those competent user judgements that concern the correctness of content from those that concern, for example, pragmatic notions of correctness, prudential correctness, politeness, etc. Anyone with knowledge of recent natural language semantics knows that this is far from trivial.

I intend to address the fundamental issues raised in this footnote in a separate publication.

is why contents are treated as truth bearers: contents are representations that are true in so far as the world is as they represent it as being. An assertion or belief will then be true to the extent to which its content is true. And truth seems to be an important dimension of assessment for assertions and beliefs.

Nevertheless, I would like to retreat for the purposes of this paper to a more general articulation of this view. According to the more general construal, assigning contents allows one (via further assumptions) to predict the conditions under which it would be *correct* to utter a given sentence. Accordingly, semantic contents come with conditions under which it is *correct* to believe or assert them. Speaking more generally of “correctness” will allow me to circumvent questions regarding the notion of truth and its role in semantics.³ It is orthogonal to the concerns of this paper whether semantic contents are evaluable as *true* in some substantive or independent sense. Those who accept the conclusions of this paper may still coherently ask themselves whether correctness as applied to beliefs and assertions may be identified with truth.

With this neutralising proviso in place, I would like to assume without further discussion the basic framework for semantic theorising that I outlined: (i) semantic theories assign propositional contents to sentences in context, (ii) these assignments ultimately yield predictions of the conditions under which it would be correct to utter a given sentence, and (iii) the contents in question are multi-purpose representational contents that can also figure as the contents of various kinds of representation, such as beliefs and assertions, and the correctness of such beliefs and assertions is at least partly determined by their contents.

For greater ease of exposition, I shall further restrict the discussion that follows to the framework of “flat” propositions or intensions. In other words, I shall be assuming further that contents are sets of evaluation points. Standardly, these evaluation points are construed as possible worlds (see, e.g. Stalnaker 1984 and Lewis 1986). Standard theories assume that the beliefs of any two thinkers who can causally interact, aim at correctness of content at the same evaluation point.⁴

However, the topic of discussion here is whether we can introduce more complex evaluation points, for example world-location pairs (“centered worlds”), and what effects this would have on the conversational model in which our contents figure. If W = the set of possible worlds, L = the set of locations, then let us say that the W^C = the set of ordered pairs $\langle w, l \rangle$, such that $w \in W$ and $l \in L$. Such centered worlds as evaluation points allow a more differentiated modelling of the conditions under which it is correct to believe a content.⁵ The contents we are operating with are then sets of centered worlds, i.e. “centered contents”. P^C = the set of such centered contents.

³ See my 2001 and 2008a for some discussion of the role of truth in semantics. I do not think the general methodological approach is weakened by speaking more generally of the conditions under which it is correct to assert or believe contents, for there will in any case be difficult problems in sorting competent user assessments properly relevant to content from other types of assessment (see previous footnote).

⁴ Stalnaker’s theory meets this condition because on his view all thinkers are actual thinkers, while Lewis meets it because on Lewis’s view, only inhabitants of the same possible world can interact causally.

⁵ I shall not, in this paper, discuss the motivation, or necessity, of operating with non-portable contents. For classic discussions, see e.g. Perry 1977 and Lewis 1979a. See also Kölbel 2013 for discussion of some motivations.

The idea is that each thinker “occupies” a certain location, and that the correctness of a belief depends on the location occupied by the believer. Thus, it may be correct for one thinker to believe a given content, while it is not correct for a different thinker to believe the same content, given that they respectively occupy different locations. In other words, there is no privileged evaluation point (such as the actual world in standard theories) which serves as the standard against which the correctness of everyone’s beliefs are measured. Rather, different world-location pairs will be relevant for assessing the correctness of different thinkers.

The centered worlds framework is a generalization of the familiar possible worlds framework. The truth-values of *some* centered contents are location-invariant (or “boring” in Egan’s 2007 terminology). In other words, some centered contents (sets of centered worlds) meet the following condition: if they contain some centered world $\langle w, l \rangle$, then for any location l^* , they also contain $\langle w, l^* \rangle$. Let us adopt the following definition:

- (P) For all $p \in P^C$: p is **portable** iff for all $w \in W$ and for all $l_1, l_2 \in L$: ($\langle w, l_1 \rangle \in p$ iff $\langle w, l_2 \rangle \in p$).

We can represent the standard framework as the system of contents whose truth-value is location-invariant, or *portable*, in this way. Thus, the standard framework is a special case of the more general centered content framework.

2. Portability and Lack of Portability

Much theorising about linguistic communication, for example Stalnaker’s and Lewis’s conversational score model (Stalnaker 1978, 2002; Lewis 1979b), uses the standard framework, i.e. assumes that all contents are portable. Simplifying somewhat, on such models information is conveyed as follows. One thinker asserts a content, a content she correctly believes, and thereby makes this content available for the assertion’s audience to believe. If the audience actually comes to believe the content thus made available, then the acquired belief will also be correct (since the source’s belief was correct, and the believed content was portable, i.e. its truth-value is location-invariant).

The idea of making available a content of belief by asserting it can be implemented in different ways. Within the (stalnakarian) conversational model I prefer, asserting a content in the course of a conversation causes the asserted content to be added to the *conversational score* (“common ground”, “presupposition set”). Given certain further linguistic or social rules, as well as the status of the conversation, contents that are part of the score become more or less suitable candidates for belief.

Whatever the specific implementation, a model involving non-portable centered contents as contents of assertion faces new questions as to how information transfer from one thinker to another can be modelled. We now have contents that it is correct to believe for some, but not all thinkers. Therefore successful information transfer can no longer generally consist in someone correctly believing, then asserting, some content, and an audience as a result coming to believe the very same content (see Stalnaker 2008, Ninan 2010a, Ninan 2010b, Torre 2010, Stojanovic 2012, Kindermann 2012, Weber 2012, Rabern forthcoming a).

It is the aim of this paper to outline a coherent way of incorporating centered contents into conversational models, specifically of treating centered contents as

contents of assertion and explaining the resulting model of information transfer. What follows will be neutral as to the specific implementation and details of a conversational model. I shall be introducing two abstract possibilities for conceptualizing information transfer with non-portable centered contents. One possibility is that a non-portable content is nevertheless *locally portable*, in which case no modification to existing models is needed. Another possibility is information transfer via a *surrogate content*: one thinker believes a non-portable content and by her assertion makes available a related surrogate content which is, in a sense to be explained, informationally equivalent.⁶ I shall sketch a communication system in which both of these possibilities are implemented and complement one another. In this system, the meaning of a sentence indicates not only the centered content the sentence has at each context of use, it also (“non-truth-conditionally” as it were) indicates the way in which that content is to figure in assertoric communication—whether updating the conversational score simply means adding the semantic content, or whether it involves adding a surrogate content.

3. Transferring information from A to B: Local Portability

The special difficulty that arises for centered contents in an account of information transfer arises because of non-portability. Whether it is correct to believe a content depends on the location “occupied by” the believer (however we cash in this metaphor). Thus, when it is correct for some source to believe a content, it may not be correct for a receiver of information to believe that content, if the receiver occupies a distinct location.

For example, suppose I correctly believe the content that my pants are on fire, a content, let us assume, that it is correct to believe only for thinkers who occupy a location at which the agent of the location is wearing pants that are on fire at the time of the location. If I express my belief by saying “My pants are on fire.”, then an audience who come to believe the same content are not guaranteed (in fact quite the contrary!) to do so correctly, because they might (and probably do) occupy a relevantly different location. Another example: suppose I believe correctly that today is Wednesday, a content that it is correct to believe, let’s suppose, only for thinkers occupying locations whose time coordinate is part of a Wednesday.⁷ I express my correct belief by writing on a postcard “Today is Wednesday.”. Then an audience who receives the postcard on a Friday (i.e. when occupying a location the time-co-ordinate of which is part of a Friday), and who comes to believe the same content immediately upon reception, is not going to be correct in doing so.

⁶ Which of the two contents should be called “the content of assertion”? In the conversational model presented here, nothing hangs on this. I have a preference for calling the content of the asserter’s belief the “assertoric content”, i.e. in general the content of the belief that would make the assertion sincere. Andy Egan in his 2007, for example, seems to prefer awarding the title to the surrogate. But I do not think that this is a substantial issue.

⁷ I am merely supposing for the sake of discussion that the content I have described as the content that today is Wednesday, and that I can express by uttering “Today is Wednesday.”, is a non-portable centered content of the sort specified. Certain views of the role of indexicals such as “today” may lead one to reject this supposition. However, the interest of this paper does not depend on the rejection of such views, for (i) there are sentences without indexicals that may be construed as expressing centered contents and (ii) this paper focuses in the abstract on communication systems involving centered contents, and how information transfer can be modelled in them. The same goes for the first example.

The basic idea of information transfer is that an informational source believes correctly and that the information contained in that belief is in some sense inherited by a receiver. If content is sometimes non-portable, information transfer cannot therefore *generally* consist in the receiver coming to believe the *same* content as the source.

However, it could do so in certain cases. The truth-value of a content can be invariant with respect to some class of locations, even if not with respect to *all* locations. Thus, even a non-portable content in the sense of (P) may be *portable relative to a class* C of locations that is a proper subclass of the class L of locations, as in (RP):

(RP) For all $p \in P^C$: p is **portable relative to a class** $C \subseteq L$ iff for all $w \in W$ and for all $l_1, l_2 \in C$: $\langle w, l_1 \rangle \in p$ iff $\langle w, l_2 \rangle \in p$.

If a source and a recipient occupy locations in a class C relative to which a content p is portable, then information transfer can be accounted for in the standard way: if the source's belief is correct and the recipient acquires a belief with the same content then the recipient's new belief is guaranteed to be correct also. We might say that a proposition is *locally portable* in a conversation just if the conversationalists occupy locations that are members of a class of locations relative to which that proposition is portable.⁸

Returning to the examples, distinct conversationalists rarely wear the same pants (and distinct pants do not necessitate each other's being on fire), so my belief of the non-portable content that my pants are on fire will only very rarely be locally portable. However, there will be many conversations in which the content that today is Wednesday will be locally portable, namely those conversations whose participants occupy locations whose time coordinate is part of the same day. Thus, in all these conversations the non-portable content that today is Wednesday could figure in the familiar way. Information can be transferred thus: a locally portable content believed by an informational source is made available through conversation to a recipient, who as a result comes to believe the same content.

4. Transferring Information from A to B: Portable Surrogates

The local portability model of information transfer only works when conversationalists occupy relevantly similar locations. However, the interest of centered contents derives in large part from situations where thinkers occupy relevantly dissimilar locations. So, unless we find an alternative model of information transfer, we will have to say that thinkers cannot transfer the information contained in their beliefs of non-portable contents to those who occupy relevantly dissimilar locations. Fortunately, an alternative model is available, and it can be integrated

⁸ If we want to allow for the locations of thinkers to change over time (in the course of a conversation), we might define *local portability* as follows:

(LP) For all $p \in P^C$: p is *locally portable* amongst persons A_1, A_2, \dots, A_n respectively at locations $L(A_1), L(A_2), \dots, L(A_n)$ iff for all $w \in W$: $\langle w, L(A_1) \rangle \in p$ iff $\langle w, L(A_2) \rangle \in p$ iff \dots $\langle w, L(A_n) \rangle \in p$.

In other words, roughly, a proposition is locally portable amongst a group of conversationalists while they are located at locations that are members of a class of locations relative to which the proposition is portable.

nicely with the local portability account.⁹ The basic idea is that a belief with non-portable content can have a portable counterpart or *surrogate*. The portable surrogate of a belief with non-portable content would be a portable content that it is correct to believe just if the belief with non-portable content is correct.

For example, if we suppose again that the content that today is Wednesday is a non-portable content which it is correct to believe at locations with a Wednesday time component, then a particular belief b_1 with that content will depend for its correctness on the location occupied by the believer. Let's call that location " $L(b_1)$ ". Then the belief b_1 will be correct just if $L(b_1)$ has a time component that is part of a Wednesday. Now, the (second-order) content that $L(b_1)$ has a time component that is part of a Wednesday would be a portable surrogate of b_1 : it contains the information encoded in belief b_1 in a portable way, for it is correct just if b_1 is correct.

In our abstract framework, we can articulate this idea as follows. First, we need to make explicit under what conditions it is correct to believe a centered content:

- (C) For all $w \in W$, for all $p \in P^C$ and all $l \in L$:
at w it is correct to believe p at l iff $\langle w, l \rangle \in p$.

This defines the relative notion of correctness at a world and at a location. But if beliefs occur at particular possible worlds, and if they (or their believing subjects) occupy a location, a particular belief can then be said to be correct absolutely just if it is correct at the world at which it occurs and at the location the believing subject occupies: We can then define "portable surrogate at a location" as follows:

- (PS) For all $p, q \in P^C$ and all $l \in L$: q is a **portable surrogate of p at l** iff:
 for all $w \in W$: it is correct at w to believe p at l iff for all $l^* \in L$, it is correct at w to believe q at l^* .

We can also define related notions of portable surrogates of particular beliefs or assertions, assuming that particular beliefs and assertions *occur at* locations:

- (C^B) For all beliefs b , for all $w \in W$: **b is correct at w** iff there is a $p \in P^C$ which is the content of b and there is an $l \in L$, such that b occurs at l and $\langle w, l \rangle \in p$.
- (C^A) For all assertions a , for all $w \in W$: **a is correct at w** iff there is a $p \in P^C$ which is the content of a and there is an $l \in L$, such that a occurs at l and $\langle w, l \rangle \in p$.
- (PS^B) For all beliefs $b, p \in P^C$ and $l \in L$: p is a **portable surrogate of b** iff: for all $w \in W$, for all $l \in L$: b is correct at w iff it is correct at w to believe p at l .
- (PS^A) For all assertions $a, p \in P^C$ and $l \in L$: p is a **portable surrogate of a** iff: for all $w \in W$, for all $l \in L$: a is correct at w iff it is correct at w to believe p at l .

In the abstract, we can simply stipulate that for each location and content, there is a portable surrogate of that content at that location. However there will be interesting decisions to be made when applying the abstract framework to concrete phenomena such as concrete episodes of thought or speech. Thus it will be an interesting question whether the semantic content of a particular English sentence can be fruitfully modelled in such a way that it is the portable surrogate of another sentence's content at some location. For current purposes, I want to demonstrate the practical

⁹ The account that follows is, though couched in quite different terms, related in spirit to Sarah Moss's account in Moss 2012.

applicability of surrogate contents as defined by (PS) by showing how minimal assumptions will guarantee that we have a suitable surrogate for a given centered content at a given location.

Let us consider an arbitrary non-portable content p , which is of course identical to a set of centered worlds $\langle w, l \rangle$ that meet a certain condition. Given that p is non-portable, this will be a condition on how w and l are related in each member $\langle w, l \rangle$ of p . If we call this relation “C” then p is the following content:

1. $p = \{\langle w, l \rangle \mid l \text{ fulfils condition C at } w\}$

Let us further consider a particular location, and let us call it “Location1”. Let us assume that this name “Location1” is portable in our meta-language in the sense that its content’s extension does not vary from one location to another.¹⁰ There will, then, be a further content q , whose members $\langle w, l \rangle$ meet the condition that Location 1 fulfils condition C at w :

2. $q = \{\langle w, l \rangle \mid \text{Location1 fulfils condition C at } w\}$

Thus, q is a portable content: whether a given centered world $\langle w, l \rangle$ is a member of q , does not in any way depend on l , but only on w (namely on whether w is such that Location1 fulfils condition C at it).

These simple assumptions concerning two contents p and q already ensure that q is a portable surrogate of p at Location1. To show this, all we need to show is that at any world w it is correct to believe p at Location1 iff for all locations l , it is correct at w to believe q at l (this follows from (PS) above). Here is the derivation:

3. For all $w \in W$: it is correct at w to believe p at Location1 iff $\langle w, \text{location1} \rangle \in p$.
[instance of (C)]
4. For all $w \in W$: it is correct to believe p at Location1 iff Location1 fulfils condition C at w .
[from 1. and 3.]
5. For all $w \in W$ and for all $l \in L$: it is correct at w to believe q at l iff $\langle w, l \rangle \in q$.
[instance of (C)]
6. For all $w \in W$ and for all $l \in L$: it is correct to believe q at l iff Location1 fulfils condition C at w .
[from 2. and 5.]
7. For all $w \in W$ and for all $l \in L$: it is correct to believe q at l iff it is correct to believe p at Location1.
[from 4. and 6.; QED]

In order to grasp a portable surrogate of an arbitrary content p at a given location l , all we need is some way of identifying l . Thus, once our conceptual resources are sufficient to grasp a content p , and once we have a (portable) grasp of a location l , we also have the conceptual resources to grasp a portable surrogate of p at l .¹¹

Another way to model information transfer within a centered contents framework is therefore this: An informational source has a belief with non-portable content while

¹⁰ The contents of subsentential expressions can be treated as portable or not in analogy to the definition of portability given above on p. 3. However, the specific way in which we would define portability for sub-propositional contents depends on how exactly we articulate the semantic properties of subsentential expressions. Since I do not want to get into this here, I shall not offer a precise definition of sub-propositional portability.

¹¹ I believe that this is an improvement over Sarah Moss’s idea (Moss 2012) that stipulations such as “I am Dr Demonstrative.” could furnish the conceptual resources to grasp surrogate contents.

occupying a certain location, and she performs a linguistic act that indicates that she has a belief with that content. Successful information transfer occurs when, as a result of witnessing the linguistic act, a recipient comes to believe a surrogate of the content believed by the source at the source's location.¹²

Returning to our two examples, the account could proceed as follows: The non-portable content I believe when I correctly believe that my pants are on fire has a portable surrogate at the location of my belief. Let's suppose that that location is the ordered pair $\langle \text{MK}, t1 \rangle$, of me and time $t1$. Then the content that MK's pants are on fire at $t1$ is a surrogate of the content that my pants are on fire at the location $\langle \text{MK}, t1 \rangle$. Similarly, the non-portable content of my belief that today is Wednesday has a portable surrogate at the location of my belief: namely the content that the day of which $t1$ is part is a Wednesday. These portable surrogates encapsulate the portable informational contents of the two beliefs with non-portable content.

5. Locally Portable Surrogates

At least two objections may be made immediately. First, proponents of a framework of purely portable contents (e.g. Frege, Perry, and many others) will object that something like the surrogate contents I am mentioning can and should figure as the contents of beliefs (and assertions) in the first place. Then we would not need a theory of portable surrogates, and could simply stick to the familiar model in which information transfer is simply content transfer. Therefore, a theory of non-portable contents that postulates informationally equivalent portable surrogates must be inferior to a theory that avoids non-portable contents in the first place. —Not so. Centered contents were introduced to do a certain job, and they still need to do that job, even if, when explaining information transfer, we resort to portable surrogates.¹³

Secondly, it might be objected that the surrogate content model does not explain the right kind of information transfer. If I know that there are certain necessary and sufficient conditions under which some belief is correct, then I can infer from that belief's correctness that the necessary and sufficient conditions for the belief's correctness are met. I can make that inference, in these very terms, without any more detailed knowledge of those conditions, but unless I understand the conditions in a way that interacts productively with other beliefs I have, and in particular with my empirical means of belief-formation, this knowledge will be trivial and useless. Knowing merely that the conditions of correctness of someone's belief are met, without any other way of conceptualizing these conditions, does not take us very far.

¹² Of course, there could also be non-portable surrogates: thinker A's belief with content p may have a non-portable surrogate for another thinker B: A's belief is correct iff B's belief is correct. But here I am only interested in identifying the notion of a portable surrogate in order to make use of it in the conversational model. Please note that the proposal for conversational updating that I will be making below does not depend on recipient beliefs always being portable (or locally portable) surrogates. The recipient belief may well be a non-portable surrogate of the portable surrogate. Thanks to François Recanatì for discussion.

¹³ As I said above, it is not part of this paper's aims to argue that centered contents are needed. Rather, assuming the need, it outlines a way of accounting for communication involving centered contents. Typical motivations for centered contents include the need to differentiate the belief I might sincerely express on 2/1/2013 by "My pants are on fire." from the belief I (and anyone else) might sincerely express by "MK's pants are on fire on 2/1/2013." This need does not disappear if we introduce the surrogates discussed here as the proper contents in the first place.

The information transfers we ordinarily effect through linguistic communication seem to take us much further than that.

This objection points towards an important problem: when we acquire new information from others' testimony this is not typically meta-linguistic or meta-mental information to the effect that some belief or some assertion is correct, or that the conditions for its correctness are met. The interesting case of information transfer is not typically one where I learn, for example, that the conditions for the correctness of my source's belief are met, at least not in those terms. When someone tells me that today is Wednesday, rather, I typically learn something from which I can infer, for example, that tonight there is a dinner at Jenny's, or that yesterday I had a metaphysics class.

If we had "portable concepts" (and perhaps portable names) for all locations, and if we had empirically adequate ways of telling which of these locations (under these very names and conceptualizations) are the ones occupied by our informants, then perhaps this would be less of a problem. However, this seems implausible. Let me therefore offer a more realistic and flexible account of surrogate contents, one that involves merely *locally* portable surrogates. A portable surrogate of a content at a location is one that is correct at *every* location iff the original content is correct at the original location (see (PS) above). A *relatively* portable surrogate is one that is correct at all location *within a certain subclass* of the class L of locations iff the original content is correct at the original location:

- (RPS) For all $p, q \in P^C$, for all $l \in L$ and all $C \subseteq L$: q is a **portable surrogate of p at l relative to C** iff:
for all $w \in W$, it is correct at w to believe p at l iff for all $l^* \in C$, it is correct at w to believe q at l^* .

A **locally portable surrogate** of content p at l within a conversation, then, is a surrogate that is portable relative to a class C of locations of which the locations occupied by the conversationalists are members.

Information transfer via merely locally portable surrogates provides greater modelling opportunities because conversationalists will always have the conceptual resources to identify relevant locations of utterances made in the conversation in empirically relevant ways that are also available to the other participants. Arguably, the content expressed by "the last but one utterance in this conversation", while being non-portable in general, is guaranteed to be locally portable in all normally functioning conversations (think of the content of subsentential expressions as portable/non-portable depending on whether they are intensions whose extensions are invariant/variable with respect to location).

Information transfer, then, when not brought about via the assertion of a locally portable content, can be effected by the assertion of a non-portable content that has a locally portable surrogate. Thus, if I assert the content that my pants are on fire, which we are supposing is the set of centered worlds $\langle w, \langle a, t \rangle \rangle$ such that a has his/her pants on fire at t in w , then even when the audience has no further identifying knowledge of me, they can, *qua* participants of the conversation, identify the following locally portable surrogate: the content that the most recent contributor to this conversation has his/her pants on fire—i.e. the set of centered worlds $\langle w, \langle a, t \rangle \rangle$ such that at t , a is at a point in a conversation where the most recent contributor has his/her pants on fire at t in w .

6. Routine Updating

Before discussing how centered contents figure in conversational models, it will be useful to explain some more general features of centered contents that flow from the existence of non-portability. Classic portable contents, as the metaphorical term “portable” suggests, can be taken from one location to another without losing their status as correct or incorrect. We can fill our memories and libraries with correct portable contents, they will not lose their correctness, no matter who believes them at what time. This is the strength and the weakness of portable contents. For even all the information in the Library of All Correct Portable Contents will not tell Dr Lauben where he is or who he is. Once we introduce extra contents to model the extra knowledge Dr Lauben acquires when he learns that he is Dr Lauben and that it is 2/1/2013, the neat portability of content is lost. For on this view, Lauben’s correct belief that he is Lauben has a content that it is not correct to believe for anyone else. Moreover, his correct belief that it is 2/1/2013 has a content that it is not correct to believe at other times.

Now, while Lauben will never stop being Lauben, presumably¹⁴ he will stop occupying a location the time component of which is on the 2/1/2013, and presumably he will at some point leave the reading room of the Library of All Correct Portable Contents. He finds out that it is 2/1/2013. A day passes, and Lauben’s experiences tell him that a day has passed. So Lauben updates his belief: he believes that it is 3/1/2013, and that yesterday was 2/1/2013. Some more time passes, but Lauben is not sure how much time has passed. So he now merely believes that it is either 3/1/2013 or later. Lauben correctly believes that he is sitting in the reading room. Then he gets up and leaves the reading room to look for the cafeteria. He will stop believing that he is in the reading room and instead believe that he is just outside the reading room in the staircase, and that he was in the reading room a moment ago. Only if Lauben updates his belief in accordance with the changes in his location will his beliefs remain correct.

To the extent to which the same thinker can occupy relevantly different locations, she must update her non-portable beliefs in accordance with these location changes in order to preserve their correctness status. A thinker might be unaware of some change in location, and as a consequence she might fail to make updates that she would have made had she been aware of the change. Similarly, she might be aware that a change has taken place, but be mistaken about the exact nature of the change. When Rip van Winkle wakes up, he is aware that his location in time has changed since he went to sleep, but he is mistaken about the amount of time that has passed since then. Unlike Rip, most of us are good at keeping track of certain routine changes and we update our beliefs accordingly. For example, if we experience an event as it happens, for example someone’s sneezing, we will at that very moment believe a non-portable content p_1 that it is correct to believe only at locations that are simultaneous with the sneeze. Just a moment later, our sense of time will allow us to realize that a little bit of time has passed. We have already stopped believing p_1 , and we have started believing a new content p_2 which it is correct to believe at locations that are a little later than the sneeze. A good while passes, and we no longer believe p_2 . We now

¹⁴ The most plausible implementations of the notion of “occupying” a location would involve thinking subjects subsequently occupying locations with subsequent time components. I am here assuming that thinking subjects persist through time. Routine updating can, I believe, be transposed into a framework in which thinking subjects are time-slices.

believe p_3 , a content that it is correct to believe only at locations that are a good while after the sneeze. And so on.

Clearly, then, a lot of updating is a matter of routine. We barely think about it. Accordingly, we don't usually classify such routine updates as changes in view, or as acquisitions of new belief. We say that in the morning, Rip van Winkle "still believes everything he ... believed"¹⁵ when he went to sleep. He didn't have time to acquire genuinely new information, or genuinely to change his mind on anything. But we don't think that this view is in conflict with the view that when about to go to sleep, Rip believes that he is about to fall asleep, and that upon waking up, he no longer believes that he is about to fall asleep. One explanation for this would be that routine updates, made as a result of routine experiences, do not count as changes in belief. Changes in belief are either changes brought about through new non-routine evidence and/or changes in portable content believed.

The same sort of routine updating will take place during any conversation. As time passes, as the conversationalists change location, centered contents believed by them will be routinely updated. This will include updating concerning the events *of* the conversation, in particular the utterances made by participants or gestures, demonstrations, grimaces made, as well as other conspicuous events taking place during the conversation. I will come back to this shortly.

7. The Conversational Role of Centered Contents

In sections 3 and 4, I outlined two basic modes of information transfer within the centered contents framework, the local portability mode and the portable surrogate mode. In section 5, I showed how both modes can be combined when information transfer is effected by transferring merely locally portable surrogates. In this section, I want to show how these basic modes can be put to use within conversational models that roughly follow the Stalnakerian and Lewisian conversational score paradigm.

The basic idea of such conversational models is that a conversation is at every stage in its course associated with a "conversational score"¹⁶ and the immediate objective of contributing to the conversation is to influence the development of the score. The conversational score corresponds, roughly, to the set of contents that are accepted by the participants for the purposes of the conversation. The score changes in accordance with certain rules of score change, i.e. rules that state how the score is to be updated when a sentence is uttered as part of the conversation.¹⁷

The best known and most discussed rules of score change are those concerning assertion and those concerning the use of sentences that involve presuppositions. I

¹⁵ Cf. Richard 1981.

¹⁶ The term "score" of a conversation is from Lewis 1979b. Alternative terms include "presupposition set" (Stalnaker 1973, 1974, 1978) or "common ground" (Stalnaker 2002). I am here abstracting from the differences between these accounts. For a detailed discussion of both Lewis and Stalnaker, see Kölbel 2011a.

¹⁷ As Lewis notices in his 1979b, if we think of *acceptance for the purposes of a conversation* as a propositional attitude of individual participants that is constitutively independent of their linguistic behaviour, and if we further assume that the conversational score is what the participants accept for the purposes of the conversation, then we cannot also say that the conversational score is constitutively governed by the rules of score change. I advocate insisting on a constitutive link between acceptance and linguistic behaviour, and not to think of it as an independently constituted propositional attitude, in Kölbel 2011a.

shall focus mostly on assertion. Most theorists, including Stalnaker himself (1978), introduce a rule of score change that says that when a content is asserted then it is subsequently added to the score, unless a participant rejects the assertion and thereby prevents the score change.

- (A) If a participant asserts that p and the assertion is not rejected by any participant, then p is added to the score.

The score, then, corresponds to what the conversationalists collectively accept for the purposes of the conversation. However, what they accept for the purposes of the conversation need not coincide with what they believe. We can think of the score as the set of those possibilities that are treated as live possibilities by the conversationalists in the conversation, and these can differ from the possibilities they regard as live possibilities *in propria persona*. Under certain conditions, perhaps in certain types of conversation, the fact that a content gets added to the score constitutes a reason for a participant to believe that content (to discard possibilities incompatible with it from his considerations). For example, if the purpose of the conversation is clearly the exchange of information, and social norms prohibit incorrect assertion in that type of conversation, then the mere fact that someone has asserted a content may by itself be evidence enough for coming to believe what they have asserted.¹⁸

The conversational score framework is usually conceived of as operating with portable contents only. Thus, when an assertion is correct in the sense that its content is correct at the location occupied by the asserter (see (C^A) above), then it is correct to believe its content at any location. The introduction of non-portable contents complicates the picture, but it does not essentially change it. We can still think of the conversational score as representing those possibilities treated as live by the conversationalists. Except that now these possibilities are centered worlds: they are centered worlds that are treated as live candidates for the purposes of the conversation. Instead of fencing off the space of possible worlds compatible with what conversationalists accept, the score now fences off the set of centered possible worlds compatible with what conversationalists accept.

What are the complications? If contents are exclusively portable, we can think of the conversation as aiming at approximation to a single designated possible world, location does not matter. Let us say that a *serious* conversation aims at correctness in the sense that it would be correct to believe the score for every participant. In a serious conversation, the target is therefore the set $\langle w, l \rangle$ of centered worlds, such that $w = @$ (i.e. $\{\langle w, l \rangle \mid w = @\}$). As long as the target set is contained in the score of a serious conversation, any conversationalist can simply adjust her beliefs to the score, and she will not go wrong.

If we allow non-portable contents, then the target of a serious conversation is smaller. To see this, consider first the target of each individual thinker in forming their beliefs. A believer occupying some location $l1$ aims at a belief set (set of centered worlds compatible with what they believe) that contains “her own” centered world, i.e. $\langle @, l1 \rangle$. As long as a thinker’s belief set contains her own centered world, her beliefs will be correct.

In a serious conversation that involves several conversationalists, each occupying a different location, the target set is greater than each individual’s, but smaller than it

¹⁸ See Kölbel 2010 and 2011a for a more detailed exposition of such a framework.

would be if all contents were portable (i.e. a proper subset of $\{ \langle w, l \rangle \mid w=@ \}$). Let us say, again, that a serious conversation is one that aims at correctness in the sense that if any one of the participants believed the score, they would be correct. Then the target subset of $\{ \langle w, l \rangle \mid w=@ \}$ is the set of centered worlds $\langle @, l \rangle$ such that l is the location of one of the participants. In other words, the aim of a serious conversation is to narrow down the score in such a way that it will contain the centered world of each conversationalist. Or more formally: the aim of a serious conversation with a set of participants P is to develop the score in such a way that it is a super set of the following target: $\{ \langle w, l \rangle \mid w=@ \ \& \ l = \text{the location of some member of } P \}$.

If, then, it is the aim of a serious conversation to develop the score in such a way that it contain the centered world of each participant, then what is the aim of non-serious conversations? One way to explain this is to say that in a non-serious conversation, the score contains unbelievably assumptions about participants' locations. Just as, in a classic portable framework conversationalists may be interested in exploring some non-actual possible worlds, in the current framework they may be interested in exploring centered worlds that are not their own. Thus, a non-serious conversation succeeds if any participant coming to believe the score would thereby have a correct belief if the centered world (or worlds) she wishes to explore were actually her own.

More generally, then, we can say that in any conversation, the participants pursue the goal of exploring some target set of centered worlds. They partially realize that goal if the score contains as a subset the target set of centered worlds. In the specific case of a serious conversation, the centered world each participant wishes to explore is their own centered world.

In the previous section, I outlined how an individual needs to adjust her non-portable beliefs in accordance with her changes in location. Some of these updates are routine in the sense that they are not perceived as a "change in mind" or as acquisition of "new" information. The score of a conversation needs to undergo the same sort of routine updating if it is to continue to be adequate to the aim of the conversation. The target centered worlds to be explored (in the serious case: the conversationalist's own centered worlds) may shift in the course of a conversation, as the assumed locations of the participants change. But, just as in the case of individual's locations, we have routine ways of keeping track of these location changes. For example, consider the content that it is 10am, a content that it is correct to believe only at locations the time-component of which is at 10am. Let's call this content p_1 . A conversationalist may assert p_1 , and it will thereby become part of the score. However, as time passes, the score must be routinely updated. Very soon after the assertion p_1 will no longer be in the score. It will be replaced by content p_2 , a content that it is correct to believe only at locations whose time component is just after 10am, and then by a content p_3 , which it is correct to believe only at locations whose time component is a little later than just after 10am, and so on.

I suggested, just now, that asserting a content like p_1 will result in p_1 being added to the score, unless someone objects. Thus, I seemingly applied a score change rule like (A) above:

- (A) If a participant asserts that p and the assertion is not rejected by any participant, then p is added to the score.

However, in our framework that allows non-portable contents, we cannot allow this score change rule generally. It will work for non-portable contents that are locally portable, as in the case where p_1 is asserted in a conversation where all participants are simultaneously present (this excludes cases where assertions are made via a postcard that arrives much later). It will not work, however, with non-portable contents that are not locally portable: suppose Lauben asserts a content p_4 that it is correct to believe only at locations the agent component of which is Lauben. Then adding this content to the score would in most cases go against the aim of developing the score in such a way that everyone's assumed centered world remains a member of it. For in most conversations, serious or non-serious, people's assumed locations have distinct agents. In other words, only a conversation in which it is accepted that all participants are identical to Lauben would be one in which p_4 can be added to the score without this violating the aim.

We therefore need to reject (A) as a general principle. We might instead accept a restricted version:

(A*) If a participant asserts that p , p is **locally portable**, and the assertion is not rejected by any participant, then p is added to the score.

This makes for a viable and interesting system of communication. However, if we wanted to say that the English sentence "I am Lauben." is an assertoric sentence with content p_4 , then this rule of score change would be silent on how the score should be updated when this sentence is used.

I therefore want to explore a slightly more complicated system which solves this problem and also opens up further attractive applications (as I hope to exemplify in the final section below). To explain this, I need to introduce the notion of "the conversational content" or "CC" of a sentence. CC is a function that takes a sentence and a conversational context as arguments and yields a content as value. When the sentence contains no CC-modifying expression, the CC of the sentence at the context is simply the sentence's ordinary semantic content at that context. However, if the sentence contains a CC-modifying expression, then the CC of the sentence at the context is a locally portable surrogate of the ordinary semantic content, one that suits the conversational situation at the context.

Let me illustrate this with an artificial example. Suppose we have a CC-modifier-free sentence s which expresses at every context a content p_5 which it is correct to believe only at locations that have Lauben as the location's agent. In other words, p_5 is the set of centered worlds $\langle w, \langle a, t \rangle \rangle$ such that $a = \text{Lauben}$ (one might say describe this as the content that one is Lauben). Because s does not contain CC-modifiers, its CC never diverges from its semantic content, i.e. for every context c , $\text{CC}(s, c) = \text{SV}(s, c) = p_5$.

Now consider a sentence s^* , which does contain a CC-modifier. For all contexts c , $\text{SV}(s^*, c) = \text{SV}(s, c) = p_5$. However, the CC of s^* at some contexts will differ from its SV. The CC of s^* at some contexts c will be a different content p_6 which is the set of centered worlds $\langle w, \langle a, t \rangle \rangle$ such that a is the speaker at c .¹⁹ This will be so at all contexts c where the ordinary semantic content of s^* is not locally portable (i.e. in this case all contexts except those very rare ones where every participant is assumed to be Lauben).

¹⁹ For concreteness, it might help to imagine that $s = \text{"One is Lauben."}$ and $s^* = \text{"I am Lauben."}$. But see §7 below.

Adding p_5 (the CC of s at c) to the score of a conversation will have the (under normal circumstances) unwanted effect of excluding from the score all centered worlds not centered on Lauben. However, adding p_6 (the CC of s^* at c) to a conversation will not have that effect: what will be excluded are merely those centered worlds whose centers are not in a conversation whose current speaker is Lauben.

The idea is that the CC of a sentence at a context will always be a locally portable surrogate of the semantic content of the same sentence at the same context.²⁰

My (programmatic) proposal, then, is to adopt the following score change rule for assertion:

(A^{CC}) If a participant of a conversation uses an assertoric sentence s in context c , and the resulting assertion is not rejected by any participant, then $CC(s, c)$ is added to the score of that conversation.

Now, if someone were to use sentence s (which has the semantic properties specified above) in a conversation, she would thereby be proposing to add p_5 to the score. If no-one objects, this should have the unusual effect of changing the assumed locations of participants: everyone will be accepted to be Lauben. If someone were to use sentence s^* , and no-one objects, they would thereby be adding a locally portable surrogate, e.g. the content that the current speaker is Lauben (or, more precisely, a content that it is correct to believe only at locations $\langle a, l \rangle$ such that a is at t taking part in a conversation in which the speaker at t is Lauben).

My programmatic proposal, then, is that assertoric sentences can be generally used to add their CC to the score, where the CC is identical to the semantic content except when sentences contain a CC-modifying device. I believe that in the abstract this makes for a fruitful model of conversation involving non-portable contents, one in which both modes of information transfer outlined in §§ 3–5 are employed. CC-modifier-free sentences can only be used to convey information in the local portability mode, while sentences containing CC-modifiers can be used to transfer information in the locally portable surrogate mode.

8. Some Proposed Applications

I have said very little so far about how the abstract communicative system (or family of systems) here sketched could be applied to natural languages. Even though the main purpose of this paper has been to adumbrate an abstract model of communication involving centered contents (that are sometimes not portable), let me conclude by at least briefly mentioning some promising areas of application.

One interesting question in thinking about applications of the conversational model sketched will be whether there are any examples of CC-modifiers, and whether we can find any pairs of sentences which generally have the same semantic content but differ in that one is CC-modified, while the other is not (much like s and s^* in the last section). I would like to offer just two examples (further elaboration will have to wait until another paper).

In general it seems to me that the classic indexicals, such as “I”, “my”, “here”, “now”, etc, are plausible candidates for CC-modified expressions. Thus, sentences

²⁰ I am, for simplicity, ignoring the possibility of a CC-modifier only modifying some but not all aspects of the location-dependence of a content.

containing “I”, “my”, “here” or “now”, etc, would seem to be good examples of sentences that on the one hand are paradigm examples of sentences to which one might want to assign centered propositions as semantic contents, but whose conversational contribution when asserted does not seem to be the addition of such centered propositions to the conversational score. So, once we adopt a centered content framework, sentences like “I am Lauben.”, or “I am a philosopher.” would seem to be sentences that express centered contents all of whose members have, respectively, Lauben or a philosopher at their center. It seems equally clear that when these sentences are used, it is not this non-portable content that is added to the conversational score, but rather some surrogate. For example, the surrogate content that the current speaker is Lauben (is a philosopher).

But it is not immediately obvious whether there are any CC-modifier-free counterparts of such sentences. There would be nothing wrong if there were no such counterparts, but in fact I believe there are some plausible candidates. Consider the following pairs of sentences:

- (1a) It’s raining.
- (1b) It’s raining *here*.
- (2a) There is a draft.
- (2b) There is a draft *here*.
- (3a) The “living statue” at the end of Ramblas is moving.
- (3b) *Now* the “living statue” at the end of Ramblas is moving.

There are some differences between the a- and b-versions of these pairs. On one possible account, “here” and “now” in the b-versions are intensional operators that shift an intensional variability present in the semantic contents of the a-versions. Proponents of such an account could claim that at any context, the a-version express non-portable centered contents, while the b-versions express the respective portable surrogate at that context.²¹ However, let us suppose that we have been persuaded by arguments like those in Lewis (1979a) that (1b), (2b) and (3b) must always and everywhere express different contents from (1c), (2c) and (3c):

- (1c) It’s raining in Poble Sec.
- (2c) There is a draft in front of the balcony door in Laura Schmidt’s penthouse in Bochum.
- (3c) The “living statue” at the end of Ramblas is moving at 12 noon, 23 minutes and 43 second, 2/1/2013.

On this assumption, then, we may have to say that the b-versions all express the same non-portable centered contents as their a-counterparts. The question then arises of what accounts for the differences between the a- and b-versions. One good proposal would be to say that while the a- and b-versions express the same non-portable contents, they contribute different contents to the conversational score. The a-versions always contribute their non-portable semantic contents, and therefore they trigger the presupposition that they are locally portable. The b-versions, on the other hand, are CC-modified. In contexts where their non-portable content is not locally portable, their conversational contribution is a suitable locally portable surrogate.

²¹ This sort of account would fit well with Kaplan’s outlook in Kaplan 1977.

Thus, if (1b) is used in a context at which “here” picks out a place that is relevantly different from the place that would be picked out if the interlocutors used “here”, then the conversational contribution will be a surrogate, such as, for example, the content that it is raining at the location of the current speaker (which is locally portable, but of course needs to be routinely updated the moment the speaker seizes to be current). Similarly, if I am sitting near the balcony door in Laura Schmidt’s Bochum penthouse, and I have reason to think that the cold draft I am feeling is present where I am sitting but not where my interlocutors are sitting, then (2b) will be an appropriate sentence to use, for the conversational contribution will be a locally portable surrogate, and not the non-locally portable content that would be contributed had I used (2a). Finally, the third example: time-sensitive non-portable contents are often at least locally portable, as in our earlier example, the content that it is correct to believe at locations whose time-component is on a Wednesday. This content will be locally portable in any conversation where all conversationalists occupy the locations whose time-component is on the same day. But sometimes the timespans within which correctness can change are much smaller, as for example when we are watching shooting stars or trying to spot the movements of a “living statue”. If, in this context, I use (3b), the conversational contribution is likely to be the content that the “living statue” was moving at a time very close to the moment at which I pronounced the word “now”. (3a) could also be taken to express a content that it is correct to believe only during very short moments (those when the “statue” is moving). But if it is used in the context I just described, then it is likely to be interpreted (by accommodation) to be sensitive to wider time spans, i.e. to express the content that the statue occasionally moves, i.e. is moving at some moments during a longer time span.

Another example that seems to me to fit the bill relies on the view that sentences concerning matters of personal taste express non-portable semantic contents.²² Thus, one might say that “Licorice is tasty.” expresses a centered content that it is correct to believe only at locations whose agent has a sense of taste that is favourably disposed towards the taste of licorice. Or more precisely: its content is the set of centered worlds $\langle w, \langle a, t \rangle \rangle$ where at t a is disposed to respond favourable towards the taste licorice has at w . Given this account of the location-variability of the contents of sentences containing taste predicates, it is also plausible to allow intensional operators that remove or modify the location-variability in question. Thus if the the content of the sentence “Licorice is tasty.” is the set of centered worlds $\langle w, \langle a, t \rangle \rangle$ such that at t a ’s sense of taste is favourably disposed towards the taste that licorice has at w , then presumably the content of the sentence “Licorice is tasty for John.” is the set of centered worlds $\langle w, \langle a, t \rangle \rangle$ such that at t John’s sense of taste is favourably disposed towards the taste that licorice has at w . In other words, roughly, “for x , —” is an operator that shifts the location-sensitivity of “—” to the location of x .²³ Now, just as the name “John” can be used in the place of “ x ”, an indexical like “I” (or, rather, “me”) can also be used. In this case we get a sentence that retains its location-sensitivity.²⁴

If this is so, then (4a) and (4b) below would seem to constitute another example of two sentences that express the same semantic contents at the same contexts, but which nevertheless differ in the way their use contributes to a conversation. (4a) would seem

²² See, for example, Egan 2010, Kölbel 2008b, Lasersohn 2005, MacFarlane 2007, MacFarlane forthcoming, Recanati 2007, Stojanovic 2007.

²³ Cf Kölbel 2009, Cappelen and Hawthorne 2009, p. 75 and Kölbel 2011b, p. 144.

²⁴ Cf. Kölbel 2013.

always to contribute its unadulterated non-portable content, while (4b) contributes an at least locally portable surrogate:

- (4a) Licorice is tasty.
 (4b) Licorice is tasty for me.

Finally, I would like to propose to apply the current framework in an account of generic “one” (and related phenomena in languages other than English). Genericity in general and generic “one” in particular, constitutes a very rich and complex topic, so I will not be able to do it justice here. However, I would like at least to gesture in the direction of an account that could be an alternative to Moltmann’s account (see Moltmann 2006 and 2010). The point is to illustrate in one more case how the current framework might be applied.

Consider the following generic sentence:

- (J) John gets cold feet in Notre Dame Cathedral.

The sentence is generic in the sense that it describes what usually or typically happens. For example, the fact that John does not get cold feet when he is walking barefoot on hot embers in Notre Dame Cathedral is not a counterexample to (J). Similarly, John might wear special feet-warmers whenever he enters Notre Dame Cathedral precisely because he gets cold feet there. It is hard to say exactly what it takes for it to be correct, and what it takes for it to be incorrect, to utter (J). But we might say that roughly what it takes is that in all normal or typical cases where John is in NDC, his feet get cold, and that what counts as normal or typical will depend on the context—e.g. if we are discussing whether and where going barefoot can lead to cold feet, normal cases will require bare feet etc. I want to assume that generally, generic sentences of this form: “_ gets cold feet in NDC.” are correct iff in all normal or typical cases where _ is in NDC, _’s feet get cold, leaving open the exact details of the analysis of genericity.

Now consider the following related sentence involving generic “one”:

- (5a) One gets cold feet in Notre Dame Cathedral.

I agree with Moltmann (2006, 2010) that “one” in (5a) requires genericity, and that the normality and typicality conditions will be context-sensitive. However, I believe that this is exactly the genericity that is already present in (J). The difference is only that while (J)’s correctness requires that in all normal or typical cases where John is in NDC, his feet get cold, (5a)’s correctness requires that in all normal or typical cases where one is in NDC, one’s feet get cold. But what exactly is this difference?

Here is a proposal: (J) expresses the following (portable) centered content: $\{ \langle w, \langle a, t \rangle \rangle \mid \text{in all normal or typical cases where John is in NDC, his feet get cold} \}$. (5a), by contrast, expresses a different, non-portable content, namely the set of centered worlds in which in all normal or typical cases where the individual at the centre is in NDC, that individual’s feet get cold. In other words, it is the set $\{ \langle w, \langle a, t \rangle \rangle \mid \text{in all normal or typical cases where } a \text{ is in NDC, } a \text{'s feet get cold.} \}$.

One prima facie problem for this proposal is that it seems that there may be occasions when it is correct to use (5a) even for thinkers with chronically warm feet, or for amputee thinkers without any feet. The account seems to predict that it would be incorrect for such thinkers to believe the centered content proposed as the content of (5a).

I agree that there are such occasions, but reject the prediction. I believe that even on such occasions it is correct for these unusual thinkers to believe the content proposed as the semantic content of (5a), for in normal or typical cases where these unusual thinkers are in NDC, they do get cold feet, for normal or typical cases will be cases where they, counterfactually, do not have chronically warm feet, or where they still have feet. Thus, at best this objection should lead us to make more explicit that these contents are modal. A more transparent specification of the content of (5a) would therefore be: $\{ \langle w, \langle a, t \rangle \rangle \mid \text{in the worlds } w^* \text{ nearest to } w \text{ if } a \text{ is normal and typical and enters NDC at } w^*, a\text{'s feet get cold at } w^* .\}$

Suppose, then, that this is a good account of the content expressed by (5a). What should we then make of the content of (5b)?

(5b) I get cold feet in Notre Dame Cathedral.

We can give the same answer as in the other four pairings: (5b) is a CC-modified version of (5a). While the conversational contribution of (5a) would be the non-portable content I specified as the semantic content of (5a), (5b) contributes a locally portable surrogate, where necessary, namely, for example, the content: $\{ \langle w, \langle a, t \rangle \rangle \mid \text{in the worlds } w^* \text{ nearest to } w \text{ if the current speaker at } \langle a, t \rangle \text{ is normal and typical and enters NDC at } w^*, \text{ the current speaker at } \langle a, t \rangle \text{ gets cold feet at } w^* .\}$

For further illustration, we can also apply this account of “one” to our examples (4a) and (4b) above. Both are most naturally interpreted as generic, because licorice is being referred to generically. This would suggest that we can replace “me” in (2b) with “one” and get a coherent generic sentence:

(2a*) Licorice is tasty for one.

Despite the slight awkwardness of (2a*), it does seem to make a similar conversational contribution as (2a).

The examples I have given in this section are supposed to illustrate how the abstract system sketched in this paper might be applied to a natural language. They are not supposed to constitute any kind of proof that the model is “the” correct or “the best” model for English or any other natural language.²⁵

References:

- Cappelen, Herman and John Hawthorne (2009) *Relativism and Monadic Truth*, Oxford: Oxford University Press.
- Egan, Andy (2007) Epistemic Modals, Relativism and Assertion. *Philosophical Studies* 133, 1–22.
- Egan, Andy (2010) Disputing about Taste. In Ted Warfield and Richard Feldman (eds), *Disagreement*, Oxford: Oxford University Press, 247–92.

²⁵ This paper has benefitted from the comments I received from audiences at three talks: in Dubrovnik in June 2012 and in Bochum and Paris in January 2013. In particular, I would like to thank Daniel Cohnitz, François Recanati, Kevin Reuter, Stewart Shapiro, Stephan Torre and Markus Werning for their feedback. The research leading to the results here presented was supported by MICINN, Spanish Government, I+D+i programme, grants FFI2009-13436 and FFI 2012-37658, CONSOLIDER INGENIO Programme, grant CSD2009-0056, as well as the European FP7 programme, grant no. 238128.

- Kaplan, D. (1977) On Demonstratives. in J. Almog, J. Perry and H. Wettstein (eds), *Themes from Kaplan* (1989), Oxford: Oxford University Press, 481-563.
- Kindermann, Dirk (2012) "Varieties of Centering and the Common Ground". Conference presentation at the workshop *Centered Content and Communication*, June 2012 in Barcelona.
- Kölbel, Max (2001) Two Dogmas of Davidsonian Semantics. *Journal of Philosophy* 98, 613–35.
- Kölbel, Max (2008a) Truth in Semantics. *Midwest Studies in Philosophy* ??, ??–??.
- Kölbel, Max (2008b) Motivations for Relativism. In Manuel Garcia-Carpintero and Max Kölbel (eds.), *Relative Truth*, Oxford: Oxford University Press, 1–38.
- Kölbel, Max (2009) The Evidence for Relativism. *Synthese* 166, pp. 375–95.
- Kölbel, Max (2010) Literal Force: a Defence of Conventional Assertion. In Sarah Sawyer (ed), *New Waves in Philosophy of Language*, Basingstoke: Palgrave Macmillan 2010, pp. 108–37.
- Kölbel, Max (2011a) Conversational Score, Assertion and Testimony. in Herman Cappelen and Jessica Brown (eds): *Assertion: New Philosophical Essays*, Oxford: Oxford University Press, pp. 49–77.
- Kölbel, Max (2011b) *Objectivity, Relativism and Context Dependence*, Hagen: Fernuniversität Hagen 2011
- Kölbel, Max (2013) Agreement and Communication. *Erkenntnis*.
- Laserson, Peter (2005) Context Dependence, Disagreement, and Predicates of Personal Taste. *Linguistics and Philosophy* 28, 643–86.
- Lewis, David (1970) General Semantics. *Synthese* 22, 18–67.
- Lewis, David (1979a) "Attitudes De Dicto and De Se". *Philosophical Review* 88, 513–543.
- Lewis, David (1979b) Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8: 339–59, reprinted in his *Philosophical Papers*, vol. 1, Oxford: Oxford University Press, 1983.
- Lewis, David (1986) *On the Plurality of Worlds*. Oxford: Blackwell.
- MacFarlane, John (2007) Relativism and Disagreement. *Philosophical Studies* 132, 17–31.
- MacFarlane, John (forthcoming) *Assessment Sensitivity: Relative Truth and Its Applications*.
- Moltmann, Friederike (2006) Generic *One*, Arbitrary PRO, and the First Person. *Natural Language Semantics* 14, 257–81.
- Moltmann, Friederike (2010) Relative Truth and the First Person. *Philosophical Studies* 150, 187–220.
- Moss, Sarah (2012) Updating as Communication. *Philosophy and Phenomenological Research* 85, 225–48.
- Ninan, D. (2010a) Semantics and the Objects of Assertion. *Linguistics and Philosophy* 33, 355–380.

- Ninan, D. (2010b) De Se Attitudes: Ascription and Communication. *Philosophy Compass* 5, 551–67.
- Partee, Barbara (2010) Formal Semantics. In Patrick Colm Hogan (ed.), *The Cambridge Encyclopedia of the Language Sciences*, Cambridge: Cambridge University Press.
- Perry, John (1979) “The Problem of the Essential Indexical”. *Nous* 13, pp. 3–21.
- Predelli, Stefano (2006) *Contexts*, Oxford: Oxford University Press.
- Rabern, Brian (forthcoming a) “Against the Identification of Assertoric Content with Compositional Value”, *Synthese*.
- Rabern, Brian (forthcoming b) “Propositions and Multiple Indexing”, *Thought*.
- Recanati, François (2007) *Perspectival Thought*, Oxford: Oxford University Press.
- Richard, Mark (1981) Temporalism and Eternalism. *Philosophical Studies* 39, 1–13.
- Stalnaker (1973) Presupposition. *Journal of Philosophical Logic*, 2: 447–57.
- Stalnaker (1974) Pragmatic Presupposition. In M. K. Munitz and P. Unger (eds), *Semantics and Philosophy*. New York: New York University Press. Reprinted in Stalnaker (1999).
- Stalnaker, Robert (1978) Assertion. In P. Cole (ed.), *Syntax and Semantics 9: Pragmatics*, 315–22.
- Stalnaker, Robert (1984) *Inquiry*. Cambridge, MA: MIT-Press.
- Stalnaker (1999). *Context and Content*. Oxford: Oxford University Press.
- Stalnaker, Robert (2002) Common Ground. *Linguistics and Philosophy* 25, 701–721.
- Stalnaker, Robert (2008). *Our Knowledge of the Internal World*. Oxford University Press.
- Stanley, Jason (2007) Introduction. In Jason Stanley, *Language in Context: Selected Essays*. Oxford: Oxford University Press.
- Stojanovic, Isidora (2007) Talking about Taste: Disagreements, Implicit Arguments and Relative Truth. *Linguistics and Philosophy* 30, pp. 691–706.
- Stojanovic, Isidora (2012) “The Problem of De Se Assertion”. *Erkenntnis* 76, pp. 49–58.
- Torre, Stephan (2010) Centered Assertion. *Philosophical Studies* 150, 97–114.
- Weber, Clas (2012) Centered Communication. Conference presentation at the workshop *Centered Content and Communication*, June 2012 in Barcelona. Available at <http://www.clasweber.net/>.