Negotiating Generics: Using Channel Theory to Model Dialogues About Natural Regularities (extended abstract)

Sheila Glasbey

Centre for Cognitive Science
University of Edinburgh
2 Buccleuch Place
Edinburgh EH8 9LW
Scotland

Email: `srg@cogsci.ed.ac.uk`
Fax: (+44 131) 650-6626

Lawrence Cavedon

Department of Computer Science
Royal Melbourne Institute of Technology
GPO Box 2476V
Melbourne 3001
Australia

Email: `cavedon@cs.rmit.edu.au`
Fax: (+61 3) 9662-1617

December 29, 1995

Negotiating Generics: Using Channel Theory to Model Dialogues About Natural Regularities (extended abstract)

Sheila Glasbey
Centre for Cognitive Science, University of Edinburgh
2 Buccleuch Place, Edinburgh UK
srg@cogsci.ed.ac.uk

Lawrence Cavedon
Dept. of Computer Science, R.M.I.T.
LaTrobe St, Melbourne Australia
cavedon@cs.rmit.edu.au

1 Introduction

Cavedon and Glasbey (1994, 1995) have recently argued that context plays an important role in the interpretation of generic sentences. In particular, certain examples suggest that a speaker has a certain collection of individuals in mind—those to which the generic relationship is applicable—when she utters a generic sentence. Cavedon and Glasbey use *channel theory*, a recent mathematical framework of Barwise and Seligman's (Barwise and Seligman 1994; Seligman and Barwise 1993) for modelling information-flow, to provide a formal framework that accounts for the context of generics. Several features of channel theory that address particular properties of information-flow also prove to be useful in addressing certain problematic issues in the semantics of generics.

An important property of generics is that they can be used in reasoning defeasibly about the properties of individuals (Morreau 1992). In the channel theoretic model, defeasible inference only involves the "relevant" individuals—i.e. those to which the generic is considered applicable. However, in a dialogue situation there may arise the situation whereby the speaker and hearer understand different contexts for a given generic—i.e. the speaker may have one collection of individuals in mind when she utters a generic sentence, and the hearer may take the generic to involve a different collection. This could lead to the hearer reaching (default) conclusions which are unintended by the speaker.

Healey and Vogel (H&V) (1994) have recently proposed a channel theoretic model of dialogue which accounts for the sort of miscommunication we are interested in. After briefly discussing the channel theoretic model of generics, we will use H&V's account of dialogue to model the information exchange between speaker and hearer, based on a short hypothetical dialogue involving the properties of an imaginary kind of bird.

2 Generics, context and channel theory

2.1 The context-dependency of generics

A generic sentence such as:

(1) Swans are white

appears to convey the existence of a "reliable" regularity—one which, roughly speaking, is "usually" or "normally" true. Such regularities have exceptions, however—we will occasionally encounter a swan that is not white—and one of the problems of giving an adequate semantics for generic sentences has been in allowing the existence of exceptions to the regularity which, nevertheless, do not undermine the truth of (1). This leads to the question of what proportion of swans need to be white in order for (1) to be true. It seems impossible to put a figure on it, and even if we managed to agree on one for this example, the figure would no doubt be different if we changed the example. Consider:

(2) Turtles live long lives[1]

which is arguably true even though most turtles are killed in their youth by predators. And some regularities appear to hold even though there may be no actual examples of them. For example:

(3) John shoots burglars

may well be true, in virtue of the fact that John keeps a gun beside his bed with the intention of shooting burglars, even though such a situation has never actually arisen.

The context-dependence of generics is highlighted when we consider pairs of examples, such as the following:

(4) Peacocks have brightly coloured tail feathers

and:

(5) Peacocks lay eggs

(4) states a characteristic of male peacocks only while (5) states a characteristic that holds only of females.[2] Yet these generic sentence seem acceptable, particularly in response to a question of the form "Tell me what you know about peacocks". Notice, however, that (4) cannot be used to argue that a given bird b is almost certainly not a peacock, given that peacocks have brightly coloured tail feathers (let us abbreviate this property to $BCTF$) and bird b lacks this property. The problem is, of course, that b may be a female. Hence, we need to be able to model the fact that (4) is restricted to male peacocks and (5) to female peacocks if we are to give a full semantic account of these generic sentences.

2.2 Classifications and channels

Barwise and Seligman's channel theory (Barwise and Seligman 1994; Seligman and Barwise 1993) is a formal framework that seems particularly well-suited to modelling generics. In particular, it provides features that allow us to account for different sorts of "exceptions" as well as accounting for the context of the law-like dependencies (or *regularities*) which a generic describes.

At the heart of channel theory is a relativistic theory of information, based on the concept of a *classification.* A classification is a structure that carves up (part of) the world into *tokens* and assigns various *types* to these tokens.

---

[1] Based on an example in (Pelletier and Carlson 1995).

[2] These sorts of examples are particularly problematic for *normative* accounts of generics, i.e. accounts which take a generic to assert a truth of "normal" individuals of the antecedent kind (e.g. (Morreau 1992)). By such an account, there are *no* normal peacocks.

Definition *A* classification *A is a structure* $\langle tok(A), typ(A), :_A \rangle$ *consisting of a set of* tokens $tok(A)$, *a set of* types $typ(A)$ *and a relation* $:_A$ *on* $tok(A) \times typ(A)$. *For* $t \in tok(A)$ *and* $\phi \in typ(A)$, *we say* t *is* classified by ϕ in A *if* $t :_A \phi$ *holds.*

The tokens of a classification may be objects, individuals, situations, or even other classifications, while the types are properties appropriate for classifying (along some dimension) these tokens. A classification can be seen as a representation of information of a particular sort. This representation is highly relativistic—intuitively, one could view a classification as being relativised to a specific agent, epistemic state, point of view, or perspective on the world.

A channel $\mathcal{C}$ is an informational link between two classifications A and B, licensing the flow of information of a particular sort. Formally, a channel is itself a classification—the tokens of $\mathcal{C}$ are *connections,* each linking tokens of A and B.

Definition *Let* A *and* B *be classifications. A channel* $\mathcal{C} : A \Rightarrow B$ *linking* A *and* B *is a classification* $\langle tok(\mathcal{C}), typ(\mathcal{C}), : \rangle$. *The tokens of* $\mathcal{C}$ *are* connections, *denoted* $s \mapsto s'$, *with* $s \in tok(A)$ *and* $s' \in tok(B)$. *The types of* $\mathcal{C}$ *are* constraints, *denoted* $\phi \rightarrow \psi$, *with* $\phi \in typ(A)$ *and* $\psi \in typ(B)$.

If $t \mapsto t' : \phi \rightarrow \psi$ holds in $\mathcal{C}$, this signifies that t being of type ϕ carries the information that t' is of type ψ. The type-token distinction allows for characterisation of different sorts of exceptions to the general rule. If a connection $t \mapsto t'$ fails to be classified by a constraint $\phi \rightarrow \psi$ and $t : \phi$, then $t \mapsto t'$ is said to be an *exception* to that constraint. Exceptions may lead to erroneous inference (i.e. a conclusion that $(t' : \psi)$ may be reached whereas the channel does not support this). Alternatively, a token t may be exceptional with respect to a constraint simply because the constraint is not applicable to it—this corresponds to the case whereby no connection in the pertinent channel emanates from t.

2.3 Modelling generics

A simple channel theoretic model of generics is obtained by using the type-level of a classification to model properties and the token-level to model individuals. A type-level constraint is used to model a generic relationship, while a connection between a token t and itself[3] is included in a channel if t is appropriate to the generic relationship represented by any constraints in that channel. For example, "Swans are white" describes a channel $\mathcal{C}$ containing connections of the form $s \mapsto s$ and a constraint $SWAN \rightarrow WHITE$, where s is taken from some collection of swans, $SWAN$ is a type that holds of swan-objects and $WHITE$ is a type that holds of white-objects. Under this model, we would expect that Australian (black) swans are not involved in tokens contained in $\mathcal{C}$ since this generic is not applicable to them. Other swan-tokens (i.e. the ones that are "normally" white) may correspond to exceptions in $\mathcal{C}$. The channel theoretic account of generics has several properties that address important issues in the semantics of generics (e.g. see (Cavedon and Glasbey 1994)).

In the channel theoretic model of generics, the content of a generic statement has both a *demonstrative content*, which is a channel, and a *descriptive content*, which is the regularity (i.e. constraint) supported by that channel. Basically, the content of a generic is the assertion that a given channel $\mathcal{C}$ supports a certain constraint $\phi \rightarrow \psi$—i.e. the assertion that $\mathcal{C}$ is of type $\phi \rightarrow \psi$.[4] The collection of tokens associated with a channel effectively captures the "scope" over which the associated generic is asserted to hold. Sentences (4) and (5) involve different channels—the first channel supports a regularity involving female peacock-tokens while the second supports a regularity involving male peacock-tokens. The issue then arises as to which channel is the "correct" (or most appropriate) one for a given generic. Basically, the appropriate channel is the one which the speaker has in mind when she asserts the

---

[3] In general, of course, a connection may involve two different tokens, but for simplicity here we assume that we are only dealing with "reflexive" connections.

[4] By allowing classifications whose tokens are channels and types are constraints we allow a uniform treatment of nested generics.

generic sentence.

The problem that is encountered in a dialogue situation is that a speaker may not impart enough information for the hearer to pick out the correct channel—i.e. the one being described. As such, the channel that the hearer takes the generic to be describing may contain more or fewer connections than the speaker intended. This may lead the hearer to conclusions that were not intended—for example, on being told that "Dutch make good sailors", the first-officer of a ship (in charge of recruitment) may sign up a lot of Dutch farmers, whereas the captain only intended the generic to apply to Dutch sailors.

This scenario—in which speaker and hearer discriminate different demonstrative contents—is addressed by H&V's channel theoretic model of dialogue. In a dialogue situation, a speaker and hearer "negotiate" until they reach (in their eyes) a consensus—i.e. if the information that a hearer is told is ambiguous or inconsistent with his current model, then he may have to clarify that information or revise his current model. H&V examine dialogues involving a maze task in which participants tend to classify the maze differently (i.e. one may use coordinates of nodes while another uses coordinates of paths between nodes). We present below a short dialogue involving the negotiation of the demonstrative content of generics and show how the dialogue can be modelled in H&V's framework.

3 A sample dialogue

In order to investigate more closely the misunderstandings that may arise when speaker and hearer differ as to which channel they regard as the demonstrative content of a generic, we examine the dialogue below. The dialogue is between a teacher (T) and a student (S), where T has given S an assignment to find out as much as possible about a species of bird unfamiliar to them both—the Tillermid.[5]

T(1) : *Tell me what you've found out so far about Tillermids.*

S(1) : *Right. Well—Tillermids are black, with scarlet wing-tips. They lay pale blue eggs and they catch fieldmice to feed to their young.*

T(2) : *OK. Tell me, are they all black with scarlet wing-tips?*

S(2) : *Well, no. Actually just the males. The females are a sort of dingy brown all over.*

T(3) : *Hmm. Bit misleading, then, isn't it, to say that Tillermids are black with scarlet wing-tips. You really must be careful to specify whether you are talking about the whole species, or only part of it.*

S(3) : *Sorry.*

T(4) : *OK, let's go on. What was the second thing? You say Tillermids lay pale blue eggs?*

S(4) : *Well, about half of them do.*

T(5) : *That's very odd. What colour eggs do the other half lay?*

S(5) : *They don't lay any eggs at all, because they're males. Only the females lay eggs.*

T(6) : *Do you think I don't know that? Don't be impertinent!*

S(6) : *Sorry.*

---

[5] We have chosen an imaginary species in order to avoid any interference with the knowledge of the reader.

T(7) : OK, so you've told me so far that Tillermids are black with scarlet wing-tips—the males, that is. They lay pale blue eggs. What else did you say?

S(7) : They catch fieldmice to feed to their young.

T(8) : Ah yes. Now, is that all of them, or just the males, or just the females?

S(8) : I don't know.

Here we have a dialogue in which the participants make a number of generic statements about the properties of Tillermids. T is requesting information from S, and there are several misunderstandings before what we might call successful communication is achieved. These misunderstandings arise when the hearer is unclear as to the intended channel described by the speaker's uttered generic sentence.

The dialogue framework below allows us to model the information exchange between T and S and illustrates the communication problems that arise in the dialogue. Note that we make no claims as to how T (the hearer) would actually revise or update her model—we do not explore this issue any further here but consider it an interesting avenue for future work.

4 The Healey/Vogel channel theoretic model of dialogue

In this section, we summarise H&V's channel theoretic model of dialogue. We only sketch their model here—see (Healey and Vogel 1994) for further details—and also modify their notation so that it is consistent with that which we use above.

H&V present a semantic framework for modelling the information flow between participants in a dialogue, which takes seriously the idea that it is unlikely that any two speakers will ever understand "exactly the same thing" by an utterance, even in cases where there is no *perceived* ambiguity or misunderstanding, and communication *appears* to be successful. H&V consider that an important source of this variation is the fact that it is unlikely that any two individuals will have exactly the same ontology—they will almost certainly "carve up the world in different ways". This, as H&V point out, raises problems for traditional semantic theories which assume an interpretation function mapping the meaning of words onto parts of the world. H&V argue that "an adequate account of multi-agent communication must address the ways in which ontological differences between interlocutors are managed". They present a formal framework, based on channel theory, which allows the explicit representation of differences in ontology between individuals.

Basically, H&V model the ontology of an agent by a classification.[6] All classifications are thus agent-specific and all information available to an agent is represented via classifications. A collection of classifications together with a number of channels is known as a *conceptualisation*. A conceptualisation is intended to model the intuitive notion of a *perspective*, such that an agent can adopt more than one perspective on the same problem. A classification is related to an utterance by a particular channel—called the *production channel*—if that utterance describes some aspect of that classification (i.e. the information that some token of that classification is classified by some type of that classification). In Figure 1, s_i and T_i are respectively a token and type from the speaker agent's classification and s_u and T_u are respectively the utterance and a classification of that utterance. Basically, the utterance is an attempt by the speaker to communicate the information that $s_i : T_i$. A hearer's conceptualisation involves a channel in which utterance tokens/types convey information relating to the hearer's classification within the hearer's current perspective. Figure 1 illustrates H&V's model of information-flow between speaker and hearer.[7]

---

[6] Actually, the ontology may be modelled across a number of classifications, each concerned with different

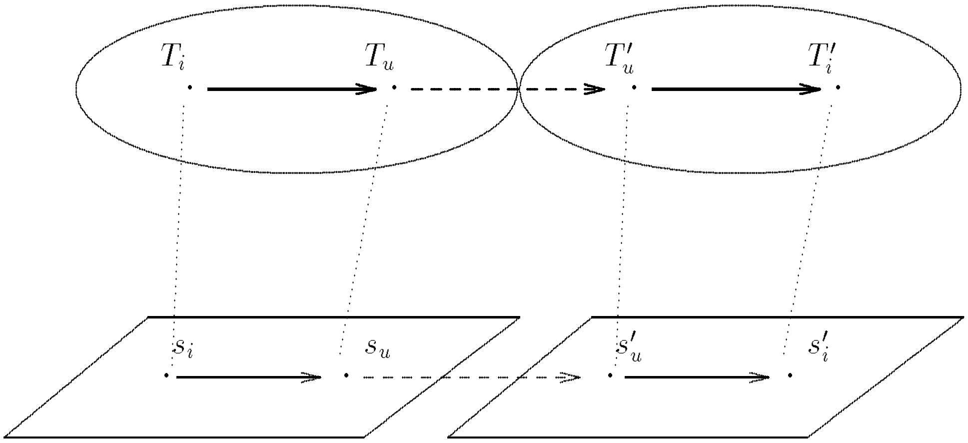

Figure 1: Channel theoretic model of communication.

As indicated in the figure, there is more than one possible way in which the hearer can interpret the utterance, either within his current perspective or by adopting a new perspective; indeed, there may be no way in which the hearer can interpret the utterance.[8] H&V classify different sorts of possible communication scenarios, based on their model:

- *clear signal:* concept token in the speaker signals a unique concept token in the hearer (via the composition of the three channels shown in the figure), under the current conceptualisations;
- *pseudosignal:* concept token in the speaker does not signal any concept token in the hearer;
- *strong multisignal:* concept token in the speaker signals multiple concept tokens in the hearer, under the current conceptualisations;
- *weak multisignal:* concept token in the speaker signals multiple concept tokens in the hearer, but under different conceptualisations.[9]

Basically, a clear signal relates to the case where the hearer *thinks* that he understands what the speaker was trying to communicate; a pseudosignal corresponds to where the hearer cannot even begin to understand what the speaker says (e.g. gibberish or mumbling); a strong multisignal corresponds to the hearer being undecided between two or more possible interpretations of the speaker's utterance; and a weak multisignal corresponds to the case whereby the hearer could render the speaker's utterance ambiguous (to the hearer) if the hearer was prepared to adopt a different view of the world.

An important property of the model is that successful communication does not necessarily occur when a speaker's concept token is a clear signal for a hearer—this case represents the scenario in which the hearer can unambiguously (within his current conceptualisation) interpret the speaker's utterance but does not necessarily mean that speaker and hearer have the "same" piece of information. However, the dialogue may proceed after a clear signal, and the discrepancy between speaker and hearer may

---

classification dimensions.

[7] We have modified H&V's model slightly to incorporate an explicit *communication channel* (discriminated by the observer) between the speaker and hearer's utterance classifications, rather than having speaker and hearer share the utterance situation. The communication channel is indicated by the dashed lines. Separating utterance situations/classifications in this way seems to better fit H&V's stance that ontologies are agent-specific.

[8] H&V also allow ambiguity between the speaker's concept and her utterance. This corresponds to the case where the speaker could have used the same utterance to describe another of her (internal) concepts.

[9] Note that a weak multisignal may be a clear signal.

only be identified at a later point. (Further details/discussion of the H&V model will appear in the full paper; see also (Healey and Vogel 1994).)

5 Applying H&V's model to our sample dialogue

In our model, the concepts or "ideas" that are communicated between the participants are the contents of generic sentences such as (4), repeated here as (6):

(6) Peacocks have brightly coloured tail feathers

We can divide this content into demonstrative content ("which channel the speaker is talking about") and descriptive content ("what the speaker is saying about the channel")—i.e. a speaker who utters the above sentence tries to communicate the information that some channel $\mathcal{C}$ supports the constraint $peacock \rightarrow BCTF$. Of course, it is not clear exactly *which* channel the speaker has in mind—this is the possible avenue of miscommunication (i.e. the hearer understanding a different channel) that we are particularly interested in.[10]

We will say that the speaker (A) has an idea I (which she tries to communicate) consisting of two components—the channel to which she implicitly refers and a description of this channel (i.e. a constraint which classifies it), consisting of some constraint. A's idea signals an idea to B (via utterance situations/types and a communication channel), which as we have seen is not necessarily the same as A's idea. B's idea will similarly have two components—the channel which he takes A to be talking about, and a description of that channel. Neither component need be the same as the corresponding component of A's idea. B may be thinking of a different channel from A, he may be classifying it in a different way, or both. In this abstract, we focus on the former case only.

We repeat our dialogue below, annotated with a discussion of how we model it.

T(1): Tell me what you've found out so far about Tillermids.

S(1): Right. Well—Tillermids are black, with scarlet wing-tips.

Here, S(1) describes some channel (call it C_{S1}) as supporting the constraint $TILL \rightarrow BSWT$ (where $TILL$ is the type of Tillermid-objects and $BSWT$ is the type of objects that are black with scarlet wing-tips), but he does not make explicit which channel he is referring to—i.e. which Tillermid token connections are contained in C_{S1}. If we say that S's idea here is I_{S1}, then we can characterise I_{S1} as $\boxed{C_{S1} : TILL \rightarrow BSWT}$. For all his hearer knows at this stage, C_{S1} could be a channel containing *all* Tillermid connections, just the male Tillermids, or just the females, for example. In H&V's terms, we can say that S's utterance here is a multisignal for T—it signals at least three "alternatives" for C_{S1}.[11]

They lay pale blue eggs

S goes on to describe some channel, which is not necessarily the same as C_{S1} as being of the type $TILL \rightarrow LPBE$. Let's call this channel $C_{S1'}$. Then $\boxed{C_{S1'} : TILL \rightarrow LPBE}$, where $LPBE$ is the type of objects that lay pale blue eggs. Once again, $C_{S1'}$ could, as far as we can tell at this point, contain all Tillermids[12], just the males or just the females. Of course, world knowledge may help us to

---

[10] There are other possible avenues for miscommunication within the H&V model, of course. For example, the speaker and hearer may have different concepts for *peacock,* or may disagree as to what constitutes *brightly coloured.*

[11] An alternative would be to think of S's utterance as vague or underspecified as to the identity of C_{S1}.

[12] Technically, we should be talking about connections involving Tillermid tokens, but this abuse makes things simpler.

decide on this—knowing that only female birds lay eggs, we would probably deduce that $C_{S1'}$ contains only female Tillermids.[13]

and they catch fieldmice to feed to their young.

Here we have a conjunction, but because it is a sentence-level (rather than a VP-level conjunction), we will assume that S may be talking about a different channel in this conjunct than in the preceding one. Again, we do not know which channel ($C_{S1''}$) S is talking about, only that he classifies it as of type $TILL \rightarrow CFFY$, where $CFFY$ is the type of objects that catch fieldmice to feed to their young. If we have no prior knowledge to help us decide, then once again we have a multisignal, signalling at least three different possible identities for $C_{S1''}$.

T(2): OK. Tell me, are they all black with scarlet wing-tips?

Clearly T is trying to ascertain here which tokens the channel contains. Let's say that the first sentence of utterance S(1) signals an idea I_{T1} to T. Then I_{T1} is $\boxed{C_{T1} : TILL \rightarrow BSWT}$. T does not know at this stage, we assume, whether C_{T1} contains all Tillermid tokens, just males, or just females.

S(2): Well, no. Actually just the males.

Thus S now reveals "which channel he is talking about"—it's one containing male Tillermids only. T can now update the characterisation of her channel C_{T1} to $\boxed{C_{T1}(\text{males only}) : TILL \rightarrow BSWT}$. (This model fits very neatly with Cavedon's (1995) method of representing implicit background assumptions of regularities via a "hierarchy" of channels. This is further discussed in the full paper.)

The females are a sort of dingy brown all over.

And from this, T now has the additional information $\boxed{C_{T2}(\text{females only}) : TILL \rightarrow DB}$, where C_{T2} is T's relevant channel in this case, and DB is the type of objects that are "a sort of dingy brown all over".

T(3): Hmm. Bit misleading, then, isn't it, to say that Tillermids are black with scarlet wing-tips. You really must be careful to specify whether you are talking about the whole species, or only some part of it.

T was obviously not satisfied by the previous answer, even though she now has the information she was seeking. T seems to be claiming here that S was wrong to offer the bald statement 'Tillermids are black with scarlet wing-tips' without making it clear which channel he was talking about. We may well feel that this objection of T's is unreasonable, given that it appears to be common practice to offer bald generic statements without making it explicit which channel is being referred to.[14] However, it may be that this is acceptable only when context and/or world knowledge makes the identity of the channel unambiguous. Alternatively, it may just not be important to the purposes of the dialogue which tokens are in the channel. Suppose that a speaker is replying to the question 'Tell me some facts about Tillermids'. Then the exact identity of the channel seems less important than if the question were 'Tell me how to recognise a Tillermid if I see one'. In the latter case, it is important to specify that the property BSWT is associated with the males only, in order to make sure that I don't fail to identify a female Tillermid as a Tillermid.

S(3): Sorry.

T(4): OK, let's go on. What was the second thing? You say Tillermids lay pale blue eggs?

S(4): Well, about half of them do.

---

[13] The processes of *how* an agent would come to such a deduction is outside the scope of the model we consider here.

[14] We even name some species after the characteristics of the male—e.g. blackbirds.

T(5): That's very odd. What colour eggs do the other half lay?

S(5): They don't lay any eggs at all, because they're males. Only the females lay eggs.

T(6): Do you think I don't know that? Don't be impertinent!

S(6): Sorry.

Here we have a misunderstanding on T's part. T(4) checks her memory of S's earlier statement. S(4), mindful of the recent reprimand he received, changes his statement that "[Tillermids] lay pale blue eggs" to the claim that "about half of them [lay pale blue eggs]". Presumably S does this because T has just told him to specify whether he is referring to the whole species or just part of it. If S were now to reply to T(4) with 'Well, the females do', we might imagine T giving him an exasperated look and reprimanding him this time for stating the obvious. As it is, S replies with S(4), and T is confused into thinking that Tillermids lay two different colours of eggs.

How does T's confusion arise? It appears to stem from the interaction between the statement 'Tillermids lay pale blue eggs', T's world knowledge and S(4). T knows that only female birds lay eggs, so she infers that the channel $C_{T1'}$ contains only female Tillermids, i.e. that $\boxed{C_{T1'}(\text{females only}) : TILL \rightarrow LPBE}$. In H&V's model, T has a clear signal when she interprets "they lay pale blue eggs"—i.e. within her conceptualisation, there is only one channel that is appropriate. Now when S(4) tells her that "about half of [Tillermids] [lay eggs]", she assumes that this refers to half the tokens in $C_{T1'}$, i.e. half of the females. This leads to an internal inconsistency for T—no channel within her current conceptualisation (one in which only females can lay eggs) can account for this statement. Since T is focussed on conceptualisations in which only female Tillermid's are considered, she interprets S(4) as leading to the idea $\boxed{C_{T1'}(\text{half the females}) : TILL \rightarrow LPBE}$. When S tells T that the half that don't lay pale blue eggs are the males, T realises that her original conceptualisation was appropriate and corrects her understanding to $\boxed{C_{T1'}(\text{females}) : TILL \rightarrow LPBE}$.

It is interesting that T(7), in going over what has already been established, says that "Tillermids are black with scarlet wing-tips", without specifying that she is referring to male Tillermids only. This is precisely what she reprimanded S for earlier. Presumably T is now depending on the fact that the nature of the relevant channel has been agreed upon and is common knowledge. This appears to happen frequently with generics—a combination of one or more factors like world knowledge, context and previous discourse often makes it clear which channel is being referred to.

T(7) . . . What else did you say?

S(7): They catch fieldmice to feed to their young.

T(8): Ah yes. Now, is that all of them, or just the males, or just the females?

S(8): I don't know.

Once again, in S(7) we have a bald generic statement with no specification of which Tillermid tokens are in the channel. And once more T(8) asks for clarification. In this case, however, S does not know the answer—S is forced to consider the possibility that alternative conceptualisations may be appropriate. That is, the appropriate channel C_{S7} described S(7) may be one containing all Tillermid tokens, just the male ones, or just the female ones.

6 Conclusion

Generic sentences seem to refer to a particular context, which is (at least in part) defined by a collection of individuals. The channel theoretic model of generics seems to capture this well, providing generic sentences with a demonstrative content. Misunderstanding in dialogue may arise when this demonstrative content is miscommunicated, potentially leading to the hearer drawing default consequences that were not intended by the speaker. (The full paper describes a particular channel theoretic model for inference with generics, demonstrating how such inferences may be drawn.)

Further work is needed to investigate in detail the interaction between information gained from an utterance and information available to the participants from world knowledge, contextual knowledge, etc. In particular, such factors seem to influence what H&V call the "current perspective". In more lengthy conversations, the role of perspective and the participants' ability to switch from one perspective to another may prove to be important. The full paper will discuss these matters further, as well as examine dialogues in which the participants differ not only in their understanding of which channel is being referred to, but may also differ in their understanding of the descriptive content (i.e. the constraint) of a generic.

References

Barwise, J. (1989). *The Situation in Logic.* CSLI Lecture Notes 17, Stanford: CSLI Publications.

Barwise, J. and J. Seligman (1994). The rights and wrongs of natural regularity. In J. Tomberlin (Ed.), *Philosophical Perspectives, Vol. 8.* California: Ridgeview.

Cavedon, L. (1995). *A Channel Theoretic Approach to Conditional Reasoning.* Ph. D. thesis, Centre for Cognitive Science, University of Edinburgh.

Cavedon, L. and S. Glasbey (1994). An information-flow model of generics (extended abstract). In *Fifth Symposium on Logic and Language*, Hungary. Revised version to appear in special issue of *Acta Linguistica.*

Cavedon, L. and S. Glasbey (1995). The role of context in the interpretation of generics. To appear in *Tenth Amsterdam Colloquium on Logic and Language.*

Healey, P. and C. Vogel (1994). A Situation Theoretic model of dialogue. In K. Jokinen (Ed.), *Pragmatics in Dialogue Management.* Gothenburg Monographs in Linguistics.

Morreau, M. (1992). *Conditionals in Philosophy and Artificial Intelligence.* Ph. D. thesis, Department of Philosophy, University of Amsterdam.

Pelletier, F. and G. Carlson (Eds.) (1995). *The Generic(s) Book.* Chicago University Press.

Seligman, J. and J. Barwise (1993). *Channel Theory: Towards a Mathematics of Imperfect Information Flow* (draft). Available by ftp as `phil.indiana.edu:/pub/SelBar93.ps`.