ELSEVIER

Cognition 53 (1994) 181–215

COGNITION

Conversation, co-ordination and convention: an empirical investigation of how groups establish linguistic conventions

Simon Garrod*, Gwyneth Doherty

Human Communication Research Centre, Department of Psychology, University of Glasgow, Glasgow G12 9YR, UK

Received 6 April 1993, final version accepted 28 February 1994

Abstract

Two experiments are reported which demonstrate the development of co-ordinated description languages in two groups of communicators playing Garrod and Anderson's (1987) maze game. The experiments contrast language co-ordination between speakers who always interact with the same partner (isolated pairs) as compared with speakers who interact with different partners drawn from the same community. Whereas the isolated pairs show higher degrees of inter-speaker convergence than the community pairs at the start of the experiment, the situation reverses by the time they have all played six or more games. The results are discussed at two levels: (1) in terms of Lewis's formal theory of conventions, and (2) in relation to a language processing model which abides by the "output/input co-ordination" principle proposed in Garrod and Anderson (1987).

1. Introduction

Most psycholinguistic accounts are based on the assumption that the language itself is given. In other words, the language, the function which relates expressions to interpretations, cannot be modified as it is being used. For most purposes this assumption is both helpful and uncontentious; but it nevertheless ignores the fact that languages are constantly changing to reflect the changing needs of the linguistic communities that support them. So there must be some psychological process embodied in communities of

* Corresponding author.

SSDI 0010-0277(94)00628-X

speakers through which a language can be modified as it is being used, and that is the topic of this paper.

Although our central concern is to examine the processing mechanisms underlying language change in a wider community, we start by considering how isolated conversationalists may exert local control over their language.[1] Given that languages are established and maintained, as well as acquired, primarily through one-to-one conversational interactions, it is reasonable to assume that the processes that underlie conversational control of the language should also have a bearing on its broader social support. So one of the key issues which we shall address is the relationship between language modification through group interaction versus modification in an isolated conversation.

First, we discuss the background to the claim that language may be refined during the course of a conversation. This process of language refinement is then considered in relation to accounts of how communities support languages as conventional systems. It will be argued that the motivation for conversational refinement of language is basically the same as that underlying convention: to promote co-ordinated understanding. However, we suggest that the kind of co-ordination mechanisms operating during isolated conversations may not be quite the same as those proposed for larger linguistic communities.

This motivates the experiments which compare the process of language refinement in isolated conversational partners with that across an experimentally simulated community of speakers.

1.1. Language co-ordination and linguistic convention

Over the last ten years there has been an increasing interest in the role of interaction in communication and its consequences for language processing (see Clark, 1985). This has been stimulated by the findings from a number of empirical studies. For instance, Clark and Wilkes-Gibbs (1986) used a version of Krauss and Weinheimer's (1964) referential communication task to examine exactly how communicators use the interaction to support reference. The first thing that they noted is that conversational references are usually established over an extended exchange. Typically, such an exchange consists of a cycle of initiation and subsequent repair only resolved when there is mutual acceptance of the referential description by both parties. So adults, like the young children in Krauss et al.'s studies, may start out with somewhat idiosyncratic descriptions but then use the subsequent interaction to repair misunderstandings, as in the following:

[1] By "language" here we mean the function relating expressions to possible interpretations mutually recognised as underlying that conversation. A more detailed characterisation of exactly what this means will emerge later in the paper.

(1) A: Uh, person putting a shoe on
B: Putting a shoe on?
A: Uh huh. Facing left. Looks like he's sitting down.
B: Okay.

In this way, dialogue partners use the exchange to build up a common interactional context against which to interpret their descriptions. That this context amounts to more than just the sum of what has been said so far was clearly demonstrated by Schober and Clark (1989). Using the same basic task they were able to show that over-hearers, exposed to exactly the same linguistic context but not able to participate in the interaction itself, perform consistently less well in identifying the intended referents.

One interpretation for these findings is that the participants are refining the language they use to suit their immediate communication goals and circumstances. This idea is central to Garrod and Anderson's (1987) account of semantic co-ordination in task oriented dialogue. They analysed a number of dialogues elicited through a computer controlled co-operative maze game where the players would regularly describe their changing position in the maze. By examining the sequences of descriptions in each dialogue, it was possible to show that players would converge on particular idiosyncratic description schemes. Each such scheme was grounded in a common conception or model of the overall maze configuration together with a restricted set of lexico-semantic interpretation rules. So, for instance, in one scheme – the line description scheme – speakers would identify a position on the maze by reference to the horizontal or vertical line of nodes on which it was located (see Fig. 1).

This type of description relies on a conception of the maze configuration which assumes that the maze is made up of horizontal rows or vertical columns related to each other in various ways. Associated with this conceptual convergence, there was also evidence of convergence on a restricted lexicon of terms to denote elements in the model. Thus a horizontal line in the scheme discussed above would be referred to exclusively as a *row*, *line*, *level*, *column*, *layer* or even *floor* depending upon the particular dialogue pair. Finally, by looking at how the various schemes were used over the course of extended dialogues, Garrod et al. were able to show that the preferred scheme for any pair could change and develop but without leading to divergence between the speakers.

So under certain circumstances, speakers are able to use the interactional opportunities presented by dialogue to refine their interpretation of the language being used. The whole process effectively generates an unambiguous language scheme or 'sub-language' associated with the task at hand which then imposes clear constraints on use and interpretation (see Garrod & Anderson, 1987 for a fuller discussion). To the extent that such semantic entrainment leads to refinement of the language, it raises the question of whether this local control through one-to-one interaction might constitute a

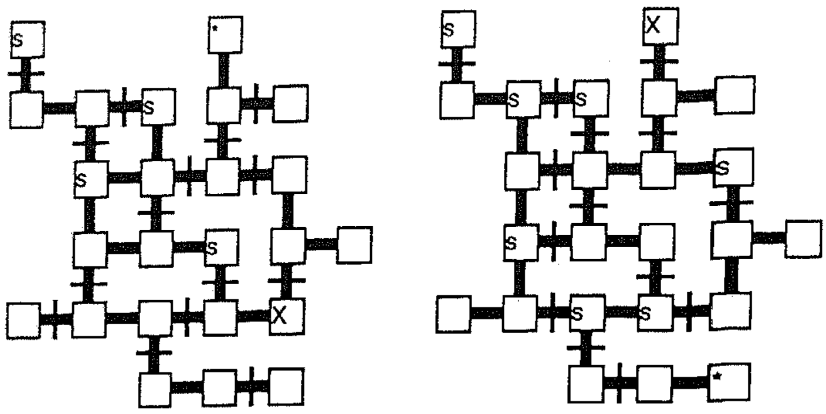

key: * = finish position, X = player's position, S = switch box, | = gate.

Fig. 1. A schematic diagram of the mazes used by Garrod and Anderson (1987). Nodes containing Ss indicate switch nodes, whereas those containing an × indicate the players' current positions and an asterisk the players' goal positions. An example of a line type of description of the point marked by the "X" in A's maze would be "On the second row up at the end on the right".

mechanism for establishing more stable conventional language systems in a larger community of speakers.

What is in common to both the process of language refinement discussed above and the process of establishing linguistic conventions is that both are motivated by problems of semantic co-ordination. Players in the maze game task adopt a particular description scheme to ensure that their interpretations are co-ordinated. In a similar vein, Lewis (1969, 1975) argues that conventions come about in situations where a *community* faces recurrent problems of co-ordination. Conformity to conventions then acts as an effective means of overcoming the problem within the community at large.

Co-ordination problems arise whenever two or more agents are confronted with choices that can only be resolved through inter-dependent decision. For instance, you are holding a telephone conversation and suddenly the connection is broken; should you attempt to call back or wait for your partner to do so? If you call back and she waits, everything is fine as it is when she calls and you wait. But if you both should attempt to call back or both wait then neither party benefits. So the root of such problems is in the inter-dependence of the appropriate choices. To succeed, each party's choice should rest on an expectation about what the other will do and it is easy to demonstrate that such reasoning soon becomes locked into an infinite spiral of higher order expectations about what your partner can reasonably expect you to expect them to do, and so on ad infinitum (Lewis, 1969).

In the absence of any secure decision mechanism, general heuristics have been proposed to explain how these problems are dealt with in practice. The most straightforward one is to choose what Schelling (1960) called the salient course of action: the one which stands out from the rest. This heuristic is effective to the extent that the choice is also likely to be salient for the other party and so tend to lead to a co-ordinated decision. Another related heuristic, appropriate for recurrent co-ordination problems, is that of precedence: choosing the previous course of action associated with that situation. So, for instance, in the telephone scenario, if one of you has regularly called the other in the past then it is best to follow the same pattern in this case. Of course, precedence is merely a source of one important kind of salience: conspicuous uniqueness of a solution to the problem because that is how co-ordination was achieved last time.

Precedence and salience seem to underlie the semantic co-ordination process operating in the maze game dialogues. Precedence is reflected in the way speakers follow what Garrod and Anderson called the "output–input co-ordination" principle, whereby in formulating an utterances the speaker will match as closely as possible the lexical, semantic and pragmatic options used to interpret the last relevant utterance from their interlocutor. Bilateral conformity to the principle quickly produces convergence on a common description scheme of the kind observed in the maze game dialogues (see Garrod & Anderson, 1987). But salience is also a factor in determining the scheme that the speakers choose to converge on. In games where they were confronted with a regular maze (as in Fig. 2), players were more likely to adopt a matrix or line type of description which takes advantage of the regularity in the maze shape, whereas with irregular mazes figural or path type descriptions were slightly more common (Anderson & Garrod, 1987). Presumably this reflects the salience of a particular scheme in relation to positions being described.

So speakers converging on local description sub-languages do so to promote semantic co-ordination. They operate according to what we shall call a *local co-ordination* process which is sensitive to the immediate constraints of salience and precedence. However, converging on common language schemes in this way does not amount to establishing a convention in Lewis's sense.

Lewis argued that a convention is more than a simple precedent and does not necessarily represent the most salient course of action. Conventions depend upon a special kind of support from the wider community. For a behaviour to become conventional in Lewis's sense it must be *common knowledge* in the community that all will conform to it on the grounds that they expect all others also to conform. In this way conventions represent rational solutions to the recurrent co-ordination problems faced by the community based on, at least an implicit, recognition that all should conform to them. Thus speakers in general are considered to operate according to a *global co-ordination* process sensitive to constraints arising from their common knowledge of that community's language system.

To this extent, it would seem that the kind of processes underlying local language refinements of the sort observed in the maze game dialogues may be somewhat different from those which underpin conventional languages of the sort discussed by Lewis. Conventions are justified because they predict what all members of the community will do in the same situation. This means that encountering one person who violates the convention is not sufficient to undermine it; there are always the others who can be expected to conform. With isolated pairs, as in the maze game experiments, this is not the case; as soon as either participant fails to support the "local convention" it loses its justification. So even though dyadic communities can support conventions, they will prove difficult to establish and maintain except as local precedents.[2] At the same time there is no evidence to suggest that the players in the maze game experiments explicitly recognised that the scheme they had adopted was common knowledge. As Garrod and Anderson demonstrated, explicit negotiation about the scheme was rare and when it did occur it had little bearing on subsequent descriptions from the pair. In cases where the players had clearly established common knowledge of a scheme through negotiation, they were as likely to shift to another scheme as in cases where no negotiation had occurred.

So to reiterate the argument, we have suggested that speakers engaged in extended dialogues can exert local control over their language relative to that dialogue. But this probably results from a local co-ordination process sensitive to precedence and salience, whereas language co-ordination in the community at large arises from a global co-ordination process sensitive to constraints on convention.

However, the Garrod and Anderson experiments only involved isolated pairs of individuals interacting over a relatively short period of time. What would happen if they were to interact over a longer period with a number of different players all drawn from the same community? Could the same kind of local co-ordination process operating within isolated pairs eventually lead to establishing a more conventional refinement of the language in the larger community? One reason for thinking that the local process may play a crucial role in more widespread language co-ordination comes from arguments about the aetiology of conventions.

Lewis suggested that conventions commonly arise in the first place out of precedents: members of the community stumble on some course of action which is a co-ordination equilibrium in that situation; this then becomes the precedent for subsequent co-ordination; the precedent becomes *common knowledge* in the community and so the stage is set for establishing a convention. Thus while conventions ultimately represent a more stable basis

[2] Lewis does give some examples of conventions among pairs of communicators but we would suggest that these are special cases which depend upon extensive shared experience of the situation.

for co-ordination than precedents they usually arise in the first place out of them.

There are also other reasons for supposing that low level processing principles such as output–input co-ordination may play a role in establishing a conventional language system in a larger community. Hurford (1989) compared a number of possible strategies that prelinguistic communities might employ to establish a co-ordinated language on the basis of one-to-one interactions. By modelling the effectiveness of various strategies, he was able to demonstrate that a so-called Saussurean strategy, similar to that of "output–input co-ordination", was optimal in establishing a common language within a community.

The experiments reported here were designed to explore this possibility by comparing the development of co-ordinated description schemes in the dialogues of isolated pairs of maze game players as compared to players drawn from a simulated community. How the community was established and the rationale behind the experiments is explained below in relation to their design.

EXPERIMENT 1: INTERACTION AS PART OF A COMMUNITY VERSUS IN ISOLATED PAIRS

The first study was designed to compare maze game dialogues from two groups of players. The first group effectively replicated the original Garrod and Anderson design except with many more games played. So each subject played nine games in succession with the same partner. However the second group interacted in such a way as to become a closed community of players. This was achieved by having each member complete the same number of games as the isolated pairs but on each occasion with a different partner chosen from the group. So over the period of the experiment, a community of speakers was created through a series of one-to-one interactions with partners drawn from the same pool. In this way, the experimental situation simulated the conditions holding in a simple linguistic community of the sort modelled by Hurford (1989).

On the basis of the introductory discussion, it is possible to make predictions about how the descriptions used by the two groups might develop depending upon which type of co-ordination process is operating. First, let us assume that the players can only co-ordinate on the basis of the local process so that any convergence on a common scheme will reflect the immediate salience and precedence of the scheme at that point in the dialogue. In this situation, we would expect to find inter-speaker co-ordination among the isolated pairs quickly approach a stable level and then stay the same for the remaining games. Though high, the level will never represent perfect co-ordination, since in any instance the salience of a particular description scheme – reflecting for instance a different type of maze on that game – could overcome the local precedent and trigger a shift in scheme. But how about the community group? If the players only have recourse to the local co-ordination mechanism, they could only be expected to achieve a limited degree of inter-speaker co-ordination on each game, since for each new partner the co-ordination process would have to start afresh. Any improvement could only arise from a general exposure to the task but without the benefits of mutually shared experience of the kind demonstrated by Schober and Clark (1989). So if all players in both groups are only operating according to the local model, there should be increasing divergence between the degree of inter-speaker co-ordination in the two groups, with the isolated pairs showing a higher level of co-ordination than the community players.

Consider now an alternative scenario where the community group begin to behave as a community in Lewis's sense. In this situation, we would expect the global co-ordination process to come into play as the experimental community becomes established. While the community group should initially be more poorly co-ordinated than the isolated pairs, we would expect the situation to change as the experiment proceeds and all converge on a common "conventional" description scheme. In fact, since convention overrides precedence and salience constraints, we would expect the pairs of players in this group to eventually become more closely co-ordinated in their use of language than the isolated pairs.

A secondary prediction relates to the relative incidence of descriptions in the two groups which only reflect a salient choice. If the community group are operating according to the global co-ordination model, then we would expect to find a smaller proportion of their descriptions reflecting local salience, since this constraint is overridden by convention.

2. Method

2.1. Subjects and design

Twenty subjects were selected from among students at a psychology summer school and each was paid £9 for participating in the experiment. They were split into two equal groups: (1) the community group, who were to play with a different partner on each game, and (2) the isolated pair group, which was split up into 5 pairs of players who then remained together throughout the study.

Each subject played nine 10 min games,[3] using a different maze configuration on each occasion. In the case of the isolated pairs, the nine games were played with the same partner throughout. In the case of the community group, the players switched partners for each new game and so by the end of

[3] This was the maximum time allowed for each game.

the study everyone in this group had played once, and only once, with each other member.

All subjects were instructed not to discuss the games with the others until the end of the experiment, which took about two weeks to complete, and, as far as we could ascertain from post experimental interviews, no significant discussion occurred.

2.2. The maze game procedure

The dialogues were all elicited through Garrod and Anderson's (1987) computer controlled maze game. Each player is seated in a different room confronted with a maze of small boxes connected by paths which is displayed on a VDU (see Fig. 1).

The purpose of the game is for the players to move their position tokens alternately through the maze (one path link at a time) until they have both reached their predetermined goals. In doing this each player can only see on the screen his or her own start position, goal position and current token position. However, they can converse freely with their partner via a high-quality audio link.

The co-operative nature of the game arises from two additional features of the mazes. First, each maze contains obstacles in the form of gates which block movement along the paths, and second, there are a small number of special boxes marked as switch positions (see Fig. 1). Both the gates and the switch positions are distributed differently across the two mazes, and it is in overcoming these obstacles that co-operation is required. If a given player (say A) moves into a box where his or her partner (B) has a switch, then all of B's open paths become gated and all the gated one's open. So when a player requires a gate to be opened they have to enlist the help of their partner, find out where he or she is located and then guide them into a switch box only visible on their own screen.

Typically a game involves players attempting to move towards their goals with dialogue intervening between moves. This dialogue contains repeated exchanges about each player's location on the maze, switch positions, goal position and so on, and it is these exchanges which we analyse.

2.3. Recording, transcribing and tagging the dialogues

The dialogues were recorded on a two-channel tape recorder, with one channel per speaker. The tape also recorded a tone produced by the computer each time the subjects made a move. This made it possible to match up the stretches of dialogue with the state of the game at that point, and so establish for each location description the exact point being described.

The dialogues were then transcribed at the word level and entered into a computer database. Those exchanges involving location description were then extracted for further analysis. A complete description exchange was defined as the cycle of conversational turns from initiation of the description through optional repair until the speakers accepted the description as complete and/or turned to another topic. Each exchange was then tagged with a code to indicate who was describing the position and the nature of the description scheme used. The basic coding system was designed to capture conceptual and semantic information about the description scheme based on Garrod and Anderson's (1987) previous analysis.

As in the earlier study most of the descriptions could be categorised according to four basic schemes, LINE, PATH, MATRIX or FIGURAL reflecting particular underlying conceptual models of the maze configuration associated with particular interpretations of the terms used in the description. A theoretical explanation of the basis for these distinctions is given in Garrod and Anderson (1987). However, to illustrate how the categorisation was done we consider below the definition of each description scheme in relation to examples of exchanges taken from the dialogue themselves.

LINE

In the LINE description scheme the speaker first identifies a line of boxes (horizontal or vertical) and then indicates the location as a box lying on this line. The following is an example of a vertical LINE description exchange from an isolated pair dialogue:

(2) B: Right, right wait a minute I'm/
A: I'm
B: First row, right my goal is the first row and the second box down. Right?
A: Right.

The main feature of LINE descriptions is that they depend upon a spatial conception of the maze as a set of rows of boxes organised in the same orientation. The rows can be labelled with a number of different terms including *level*, *column*, *line* or just *row* and can be organised either vertically (as in the example above), horizontally or even, on occasion, diagonally with respect to the maze.

PATH

In a PATH description the speaker starts by identifying a prominent point on the maze and then describes a path which follows the links and boxes ending at the location to be described. An example from an isolated pair dialogue is shown below:

(3) B: Right guide me to a switch. I'm two along from you.
A: Right.
B: Two from the right from you.

Path descriptions rely upon a very direct topological model of the maze configuration reflecting paths and/or the boxes linked by those paths. Such description schemes may differ according to the elements actually counted in the description. This means that the same location can be described in different ways depending on whether it is the boxes traversed that are counted or the path links between them. Counting boxes leads to a description of point *X* in Fig. 2 as "two along and one up" whereas counting path links leads to "one along and one up". So even this most concrete of description schemes is open to misinterpretation.

Apart from the basic PATH descriptions, a special type was also identified which we call Goal related PATH description. These tend to occur towards the end of a game when the players are negotiating how to reach their goals simultaneously. At this point they commonly indicate the position as a path distance from the goal:

(4) A: Right I'm one away from my goal.
B: Are you?
A: Uh huh provided it stays like this.

We treated these as a distinct subcategory because they commonly serve a rather different function from the other descriptions: acting as comments on the player's position relative to the state of the game rather than simple locatives. Furthermore, as we shall see below they can represent quite a large proportion of the description exchanges in some games.

MATRIX

In a MATRIX description the speaker identifies two sets of orthogonal lines on the maze, one in the horizontal plane and the other in the vertical,

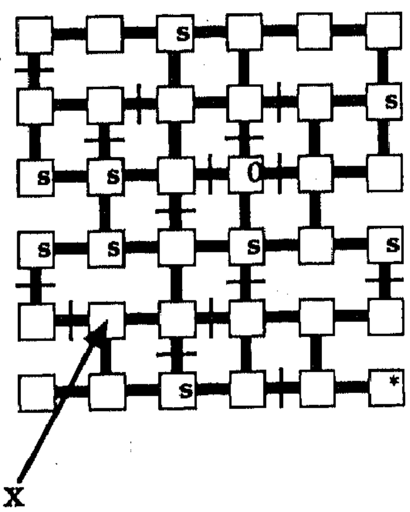

Fig. 2. A schematic diagram of a symmetrical maze, with the point X described in the text indicated with an arrow.

indicating the position as at the intersection of the two. This example comes from a dialogue between a pair from the community group:

(5) B: Mine's in E2 my wee man/
A: Right.
B: And my goal's in C4.
A: OK.

The main characteristic of such descriptions is in the use of counting codes, which may employ two sets of numbers or on occasion letters (a = 1, b = 2, etc.) in combination with numbers, as in the example given above. The codes are then used to identify each of the orthogonal lines of boxes. Although such descriptions can look very similar to path descriptions there are two clear differences. First, in defining positions the matrix scheme takes no account of the links between the boxes, it only depends upon identification of the lines containing the box to be described. Consequently, there are many locations that cannot easily be described with the PATH scheme due to missing links along the most straightforward path, but which present no such problem for MATRIX descriptions with the same origin (see Garrod & Anderson, 1987). The second difference concerns the method of identifying the intersecting lines. Two possible PATH descriptions were offered above for point *X* in Fig. 2, corresponding to "one along and one up" or "two along and one up", neither of these combinations could yield a MATRIX description of the same position since they do not give an independent description of the two lines. The comparable MATRIX description for this point would normally be "two, two" corresponding to the second line up and the second line across.

FIGURAL

In the FIGURAL type of description the speaker identifies some particular configuration or pattern of boxes on the maze and then indicates a position relative to this pattern. Here is an example where the figure is described as a "wee square of boxes":

(6) B: You don't have a lever in eh you know the wee square of boxes up in the left/
A: whoops
B: the right hand one of that one?

Here is an example where there is a reference to an 'L' shape:

(7) A: I'm not in that L shape I'm in the like the second row downwards – column downwards.

Apart from the four basic schemes described above we also encountered a few description exchanges where the description itself was only given obliquely as part of a comment about a position already known to the players. An example of such a description in relation to a switch box is shown below:

(8) B: I had a switch in there that should've worked, och well.
A: You had a switch in there so it changes me.

These we classified as comment descriptions and they were treated differently in the overall analysis.

Finally, there was a small residue of exchanges which were coded as undefined relative to the four basic schemes. In most instances, this was simply because the description was incomplete at the end of the exchange and/or aborted as in the following example:

(9) A: Right where's yours?
B: Can you move up?

or on occasion, just too vague to classify (e.g., "I'm, sort of, in the middle".) These were coded as "undefined descriptions".

All such exchanges and those classified according to the schemes described above were coded initially by the second author and then all the exchanges in a random sample of 25% of the dialogues (i.e., 574 exchanges) were independently coded by the first author. The agreement between the two coders was 93% overall.

3. Results and discussion

First we consider the overall distribution of different types of description exchanges in the isolated pair and community group dialogues to establish the degree of convergence onto common schemes. We then analyse the sequence of exchanges within each dialogue in order to establish the relative degree of inter-speaker co-ordination in the two groups and test the various hypotheses introduced above.

3.1. *Distribution of exchanges using the different schemes*

The 45 dialogues in the community and isolated pair groups contained 1377 and 1052 description exchanges respectively. These were distributed quite evenly across the games in both groups. For the initial analysis they were organised into three sets reflecting different degrees of game experience: early games, where the average experience of the two players was 3 or less games each; middle games, where the average experience was between 4 and 6, and late games where it was greater than 6. The overall distribution of the different types of description is shown in Fig. 3.

Looking first at the community group, it is apparent that the distribution changes over the three sets of games. Whereas in the early game dialogues as many as 22% of all exchanges used LINE descriptions, with about 58% using MATRIX, for the middle and late games the MATRIX type predominated representing 80% of all exchanges (including those which were undefined, comment or goal related PATH). If we ignore all descriptions that are not in the four main categories 98% in the Middle games and 100% in the late games were of the MATRIX type.

In contrast, the overall pattern for the isolated pairs is very different. From the early games onwards this group produces a substantial proportion of both LINE (27%) and MATRIX descriptions (40%) replicating the pattern observed by Garrod and Anderson (1987).

This difference between the two groups arises from the way speakers develop their preferred schemes in the initial games. Figure 4 shows a more detailed breakdown of each speaker's description choice in their first five games represented as relative proportions of the four main types. The proportions of MATRIX descriptions are also broken down into three sub types reflecting different ways of assigning letter and number codes to the two axes: $MATRIX_1$ = letters across the way and numbers down, $MATRIX_2$ = numbers across and numbers down, and $MATRIX_3$ = numbers across and letters down.

As can be seen from the figure, the isolated pairs settled down within a game onto their preferred schemes, with two pairs adopting MATRIX (either subtype 1 or 3) and the remaining three pairs using LINE. For the community group this LINE–MATRIX pattern is much the same for the first game but over the next three games MATRIX comes to dominate, with one sub type $MATRIX_1$ becoming paramount. The data from the two groups were compared by entering each subject's proportion of MATRIX descriptions per game into a two-way mixed design ANOVA with game (games 1–5) as a within subjects factor and Group (community vs. isolated pair) between subjects. The ANOVA revealed a reliable main effect of group, with $F(1, 18) = 4.5$, $p < .05$, reflecting the greater proportion of MATRIX descriptions used by community pairs versus isolated pairs, as well as a main effect of game, with $F\ (4, 72) = 9.2$, $p < .01$, reflecting the overall increase in the proportions of MATRIX across games. There was also a reliable interaction between group and game, with $F(4, 72) = 9.324$, $p < .01$, reflecting the steady increase in proportion of MATRIX descriptions for the community group as compared to the isolated pairs who did not show any increase (for simple effects of game at community group, $F(4, 72) = 18.65$, $p < .01$; for simple effect of group at games 4 and 5, $F(1, 26) > 8.9$, $p < .01$).

So it seems that the various speakers in the community group rapidly converge onto one scheme. Furthermore, the scheme they converge on is

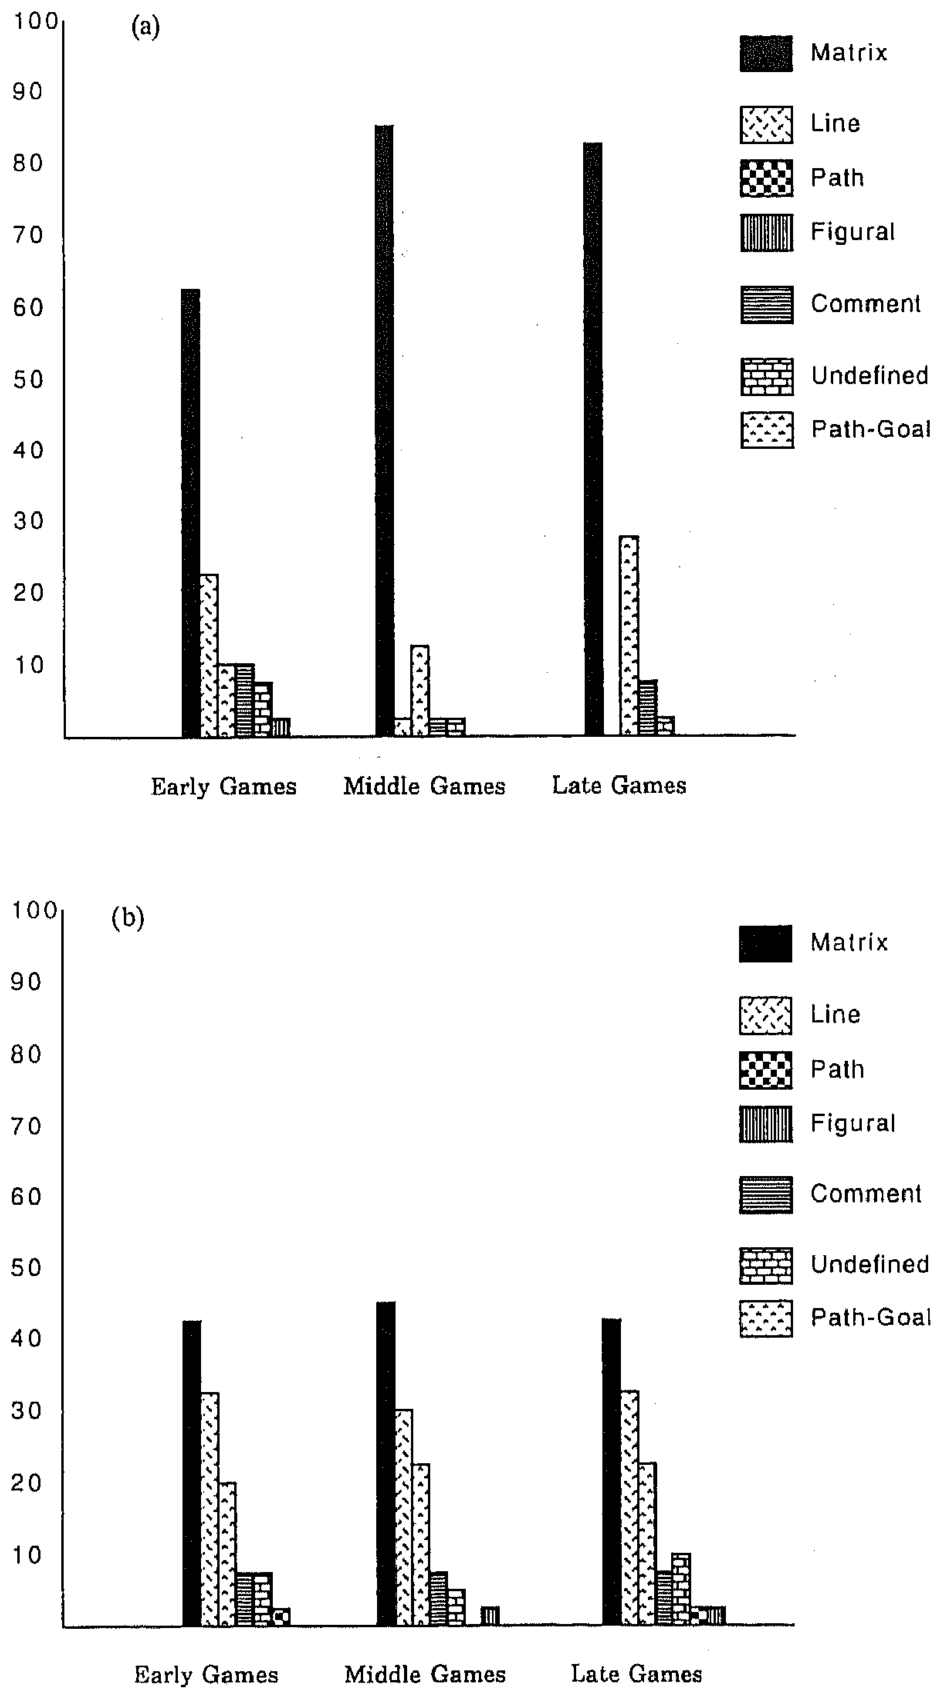

Fig. 3. Overall percentage of description exchanges falling into each description scheme type for early (1–3), middle (4–6) and late (7–9) games. (a) Data from the community group. (b) Data from the isolated pairs in Experiment 1.

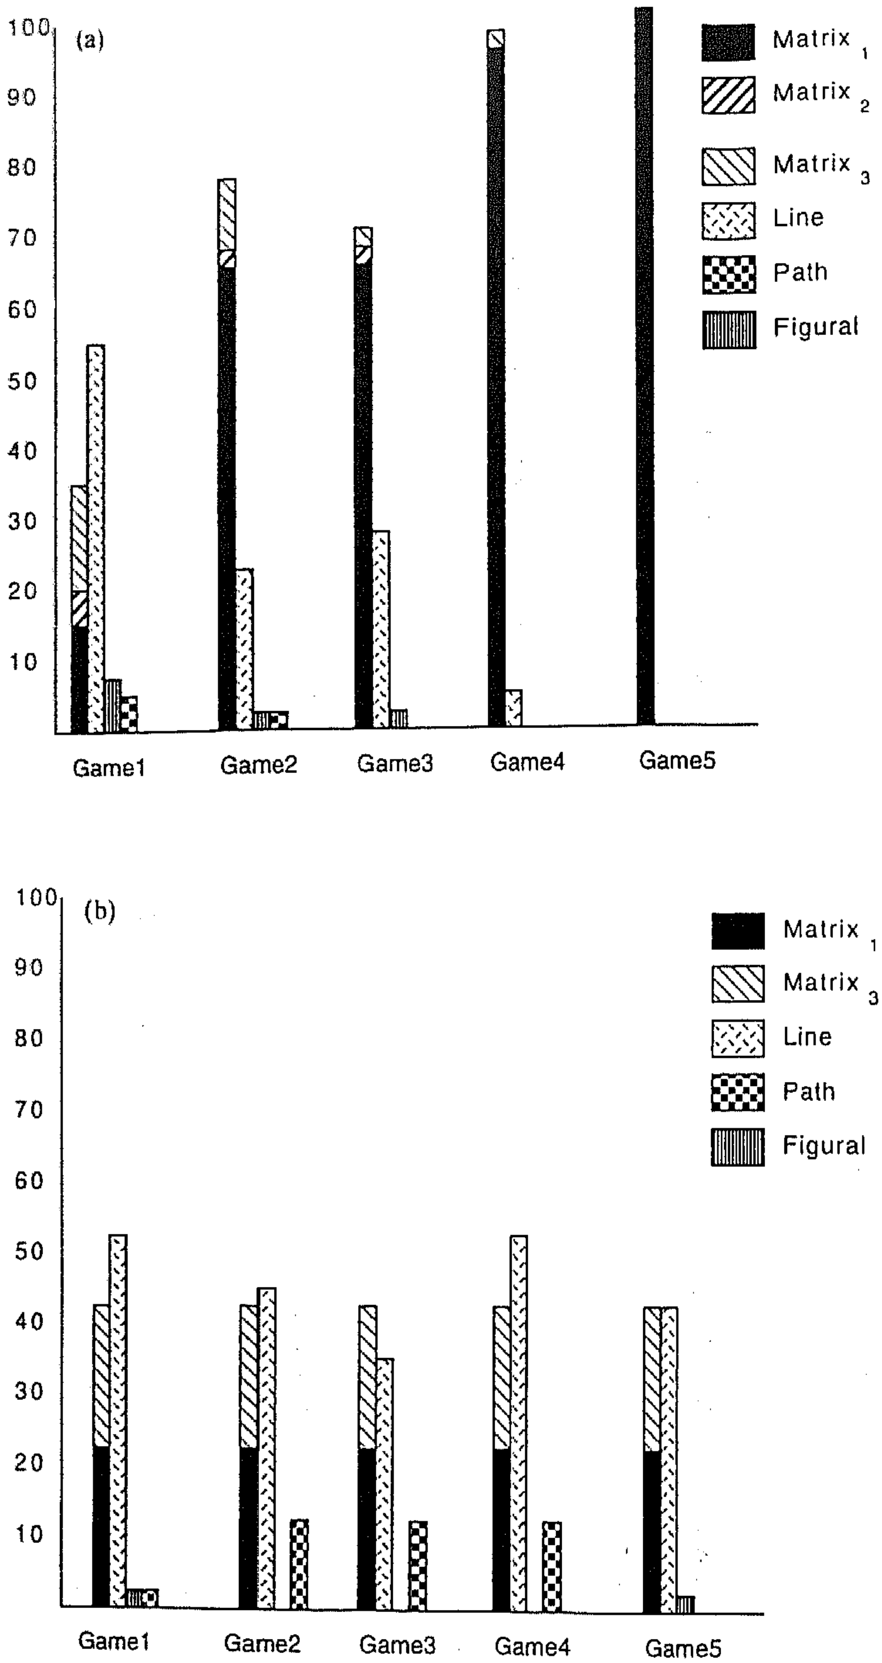

Fig. 4. Average percentage of description exchanges of each type for each player over their first 5 games. (a) Data from the community group. (b) Data from the isolated pairs in Experiment 1.

the most popular one overall in their initial games: for game 1, MATRIX represents 49.5% of all exchanges across the community group as a whole, and, within these, the most commonly used sub type is MATRIX_1, which represents 50% of all MATRIX descriptions.[4] So any convergence mechanism would seem to promote what is initially the most commonly used description across the community as a whole. Before discussing how this might come about, we consider some more direct analyses of inter-speaker co-ordination.

3.2. Measures of inter-speaker co-ordination over the course of the experiment

The main predictions from the alternative models discussed in the introduction concern how co-ordination in choice of scheme should develop across early, middle and late games for the two groups of players. To test these predictions requires a sensitive measure of language entrainment that can be applied independently to each speaker in each dialogue. The one which we use here reflects the degree to which any speaker's choice of description for exchange n in the dialogue is predicted by their partner's choice in the previous exchange n-1 (see Garrod & Clark, 1993). This was calculated by taking the sequence of description exchanges[5] in each dialogue and establishing the number of cases when A's description in exchange n matched B's description in exchange $n-1$. The result was then divided by the total number of transitions in that game (i.e. cases where the preceding description had been given by that player's partner). If choice of scheme is perfectly co-ordinated between the two speakers, this would produce a sequential matching probability score of 1 for each player. Lower scores would reflect decreasing degrees of entrainment to a chance baseline of 0.469, and can, of course, be quite different for the two players in the same dialogue.[6]

The mean scores for each group are shown in Fig. 5 pooled across early, middle and late games. Although speakers in both groups produce above

[4] Note that these figures are for the number of descriptions used across the whole group and are substantially higher than the average proportion used by each speaker. This comes about because speakers who adopt the MATRIX scheme tend to produce more description exchanges per game than speakers who adopt other schemes. Perhaps this reflects in some way the effectiveness of the scheme.

[5] These sequences were extracted just on the basis of exchanges which used one of the four standard description schemes: PATH, MATRIX, LINE, FIGURAL. So they did not include exchanges involving Goal related path descriptions or comment type descriptions since it was judged that these tended to occur in special context.

[6] The chance baseline was calculated on the basis of the relative proportions of each scheme in all the first games. The measures are independent for each speaker in a dialogue since it is possible to have a situation where A has a sequential probability of matching measure of 1 while B has 0. This would happen if A always matched B's previous descriptions and B, somewhat obtusely, always chose not to match A's previous ones.

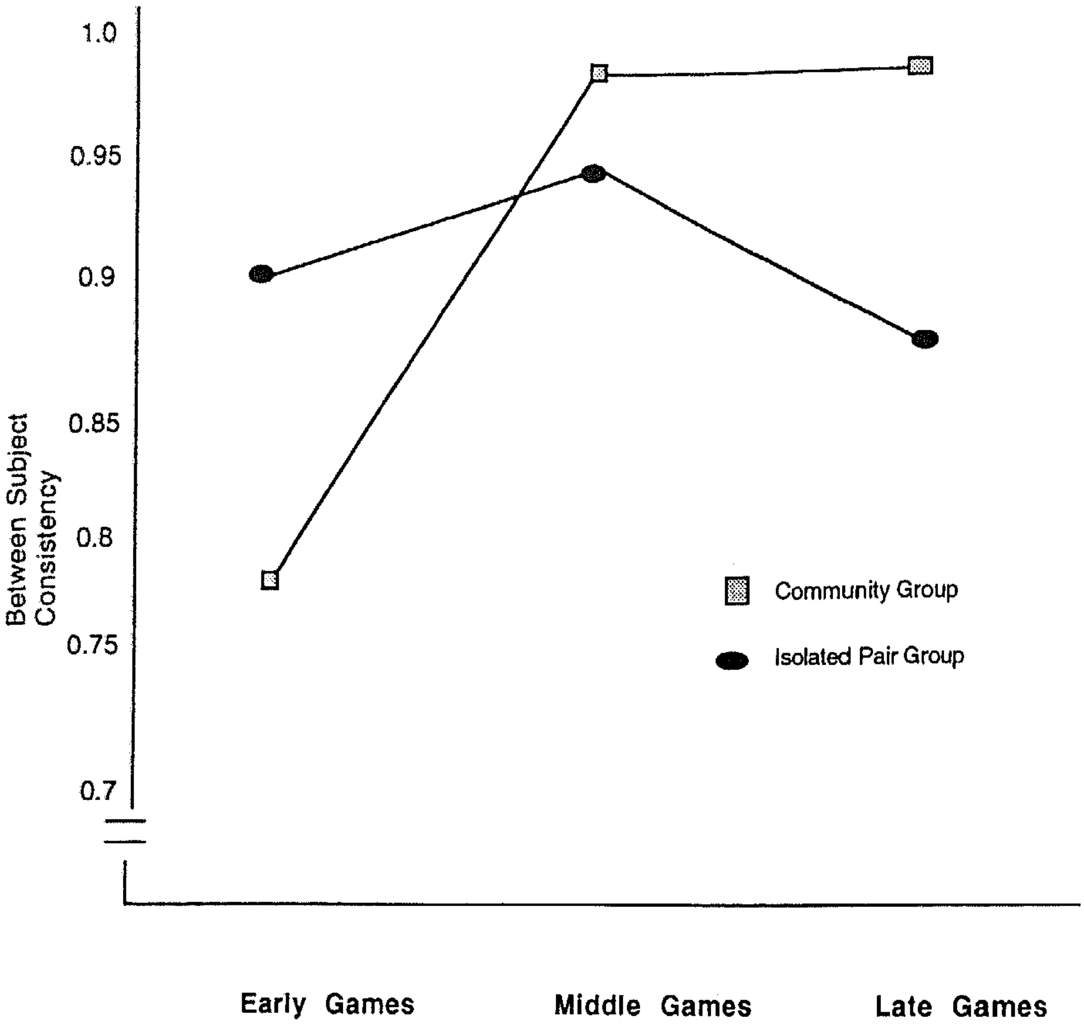

Fig. 5. The mean sequential probability of matching the previous speaker's description in sequences where the speakers change. The consistency measure is shown for early (1–3), middle (4–6) and late (7–9) games for both the community group and the isolated pairs in Experiment 1.

chance scores for all three sets of games ($p < .01$, sign test) it is apparent that the degree of inter-speaker co-ordination develops differently for the two groups. To examine this relationship an ANOVA was carried out where each subject's score was entered into a mixed design, with game (early, middle and late) as a within subjects factor and group (isolated pairs vs. community) between.

The ANOVA revealed a main effect of game, with $F(2, 36) = 7.512$, $p < .05$, as well as an interaction between game and group, with $F(2, 36) = 10.078$, $p < .001$. Further analysis of the interaction showed that the increase in co-ordination across game is reliable for the community group ($F(2, 36) = 16.316$, $p < .001$, for the simple effect of game for this group) but not for the isolated pairs ($F(2, 36) = 1.27$, $p > .1$). So the community players exhibit a steady increase in the degree of inter-speaker co-ordination as each is exposed to more members of the group, whereas the isolated pairs show no improvement over the level achieved in the first three games.

Comparing the effects of group at the various game levels also reveals that for the early games the isolated pairs are more closely co-ordinated than the community pairs ($F(1, 45) = 8.976$, $p < .01$, for the simple effect of group at early), but by the time they reach the late games, this situation has been completely reversed with the community pairs now showing the closer co-ordination ($F(1, 45) = 4.154$, $p < .05$, for the simple effect of group at late).

So these results are very much in line with the predictions of the global co-ordination model. In the early games it was predicted that the community pairs should be less co-ordinated than the isolated pairs, since they would not have had the opportunity to establish a community. But by the end, it was predicted that they should demonstrate closer inter-speaker co-ordination due to the stabilising influence of establishing a community wide convention.

One of the reasons for predicting this more stable and elevated degree of co-ordination comes from a basic contrast between the two types of co-ordination process. We argued that the global process establishes a convention in the group, making speakers relatively impervious to the local context, whereas the local process produces co-ordination as a direct consequence of sensitivity to that local context in terms of immediate salience and precedence constraints. This means that when formulating a description there will sometimes arise a conflict between following the local precedent or adopting a more salient scheme, and whenever the salient description is chosen it will reduce the degree of inter-speaker co-ordination.

If this explanation is correct, we would expect to find a much lower incidence of descriptions in the community group dialogues reflecting local salience. This hypothesis is tested below.

3.3. Differential effects of salience in determining choice of description scheme

In the context of the maze game there are certain occasions when a particular type of description becomes especially salient for the speakers. One such occasion is toward the end of the game when both players are attempting to reach their goals simultaneously. At this point, it is very natural to indicate your position in relation to the goal, with some such description as "I'm now two away from goal". For this reason, these descriptions were classified separately in the original analysis as goal-related PATH descriptions and represent about 16% of all description exchanges in the corpus. Because of their special nature, such descriptions were excluded from all the analyses of inter-speaker co-ordination described above. However, since they are nearly always chosen in conflict to the previously established scheme, they do represent a case where local salience can be seen to override precedence.

The proportion of goal-related PATH descriptions produced by each player was therefore calculated for the early, middle and late games. For the isolated pairs it increased from 22% (early) to 26% (in both middle and late), whereas for the community group, it started as 9% (early) increasing to only 12% (middle) and 15% (late). The individual subject's data were entered into an ANOVA design, in the same way as the inter-speaker co-ordination measures discussed earlier. It revealed both a main effect of group, with $F(1, 18) = 5.977$, $p < .025$, and a main effect of game, with $F(2, 36) = 3.913$, $p < .05$. So it seems that the basic prediction is upheld: the isolated pairs generate about twice as many goal-related PATH descriptions as the community players. But it is also apparent that the proportion of these descriptions tends to increase across games for both groups.

This latter finding is almost certainly due to the fact that, with more experience, the players are more likely to complete the games within the allotted time and, as a result, encounter more opportunities to use the goal related descriptions. The completion rate for the community group increases from 30% (early) to 70% (middle) and 93% (late), whereas for the isolated pairs the pattern is 67%, 80% and 93% respectively. Thus even at the end of the experiment, when the community group players are completing 93% of their games, they still use fewer goal-related PATH descriptions than isolated pairs completing 67% of their games at the beginning of the experiment.

So the evidence indicates that the community group are less sensitive to the influence of salience in their choice of description scheme; but we would also expect them to be less sensitive to the influence of local precedence if the choice of scheme is conventional. This hypothesis can also be examined in relation to the incidence of goal-related PATH descriptions.

3.4. *Differential effects of local precedence in choice of description*

As we have argued above, goal-related PATH descriptions tend to be introduced into the dialogues in conflict with the previously adopted description scheme. So their initial use gives some indication of a speaker's sensitivity to salience in choice of description. However, once one of the speakers has introduced the new description, this then becomes an immediate precedent for the next exchange. So its subsequent use could also indicate sensitivity to local precedence.

All the later game dialogues (i.e., middle and late) were therefore examined to establish the sequential probability of matching the goal-related PATH description on the next exchange. In the community group dialogues, there were 34 cases where a previous goal-related PATH description was matched in the next description exchange, as opposed to 53 cases where the speaker reverted to the previously established scheme. So this represents a conditional matching probability of 0.395 overall. However, with the isolated pairs there were 59 cases where the next description matched, as opposed to 53 where the speaker reverted to the previous scheme, yielding a conditional matching probability of 0.518 overall. Although based on a rather small amount of data, the difference in each speaker's probability of matching goal-related PATH descriptions is marginally reliably by t-test, with $t(18) = 1.483$, $0.05 < p < .1$. So it would seem that there is some indication that the community group are not only less influenced by salience, but also less influenced by local precedence.

4. Discussion of Experiment 1

The overall pattern of results from this study clearly points to a quite dramatic difference in the way the two groups of subjects co-ordinate their language use across games. The isolated pairs behave very much as expected on the basis of Garrod and Anderson's earlier study. Within a game or two each pair has achieved a reasonably high degree of inter-speaker co-ordination which does not increase over the course of the remaining games. Furthermore, their choice of description scheme seems to be constrained both by salience, as exemplified in the proportion of goal-related PATH descriptions, and local precedence, as indicated by the generally high level of within pair entrainment as well as the trend toward increased conditional probability of matching goal-related PATH descriptions.

The community group, on the other hand, perform quite differently. In the early games, the pairs are less co-ordinated in their language use, behaving as if primarily under the influence of weak local co-ordination constraints. However, as soon as the community becomes established, through overlapping interactions, they begin to co-ordinate as a group, subject to the same kind of global co-ordination constraints that apply to communities in general. This produces both an exceptionally high degree of inter-speaker co-ordination in choice of scheme and a much reduced susceptibility to using the more salient goal-related PATH descriptions towards the end of each game.

The results raise interesting questions about how these differences in degree of co-ordination actually come about. The two main findings in need of explanation are: (1) how members of the simulated community manage to converge onto the same scheme so quickly; and (2) how the community members then become so closely co-ordinated with each fresh partner.

We shall start with the first of these. What needs to be explained is how community group members can all converge onto one scheme when they have only encountered less than half of the other community members. The result is particularly striking when one considers that the subjects have no basis for even believing that they are part of such a community.[7] This means that any model of the process will have to explain the rapid convergence without recourse to assumptions of explicit knowledge of community membership. The model which we shall consider makes two basic assumptions: first, that on each new game the two players will always attempt to converge onto a common scheme; and second, that whenever there is a conflict, the scheme which was most commonly used by both players in their immediately preceding games will dominate.

First, we test how well this model fits the actual pattern of dialogue outcomes in the community group as a whole, and then go on to explore its more general implications. The model can be tested by considering all the games which occurred before complete convergence onto the MATRIX_1 sub-scheme. It was possible to identify 19 games in this category, and, in all but one, both speakers ended up sharing the same predominant scheme (i.e., the most commonly used description for each speaker in that game was of the same type). So the first assumption about game by game convergence is well supported by the data.

To test the second assumption we have to consider situations where the two players had adopted different schemes in their immediately preceding games, in other words, where there was a potential description conflict between the two speakers at the beginning of the new game. There were 9 such cases among the 19 pre-convergence games. In 8 of these, the scheme that the speakers converged on matched the most commonly used description across the two speakers' preceding games. So, if on their previous games one speaker had produced 18 MATRIX descriptions and the other 10 LINE and 2 MATRIX, or just 10 LINE alone, then MATRIX would predominate when the two came together. The only game going against this pattern was a case where the two players had previously adopted opposing MATRIX sub schemes (sub-schemes 1 vs. 3), and they both ended up converging on a LINE type description. But overall, the pattern of dominant scheme from the preceding games winning any conflict is reliable by sign test ($p < .05$) against a null hypothesis that choice of scheme is random so long as it matches either one or other partner's previously adopted scheme.

So a global co-ordination model which assumes pressure to converge with each new partner linked with dominance of the most commonly used past scheme fits well the pattern of outcomes seen in our community group. Such a model predicts rapid convergence in a closed community of this kind because each new convert goes back into the pool to convert the others, and, it promotes the most popular scheme overall because that is going to have the highest chance of winning early on and so will increasingly dominate in subsequent encounters.

[7] Recall that the players are only aware of encountering a fresh partner on each game, and even then, only encountering them remotely, since they are sitting in different rooms.

However, before we go on to consider the broader implications of this convergence account and how it might be expected to lead to increased levels of inter-speaker co-ordination, we need to make sure that there are no other factors which could also explain the pattern of results. One obvious difference between dialogues in the two groups was the incidence of explicit negotiation of the descriptions, that is cases where the speakers actually talked about the form of the descriptions to be used. As in Garrod and Anderson's original study, all such cases were identified and can be compared across the two groups.[8]

For the isolated pairs there was some negotiation in 20% of early games and 7% of middle games, but none in the late games. This contrasts with the community group where 47% of early games, 100% of middle and 87% of late games contained some explicit negotiation of the scheme. Furthermore, in the community group games, the negotiation was of the same form in nearly every case: the discussion occurred at the beginning of the game and was used to check the details of the scheme to be used. A typical example is shown in the extract below which is taken from the beginning of a middle game from the community group:

A: How do you want to name the columns?
B: Eh letters across the top.
A: From the left?
B: Hmm well no we're Japanese do it the other way and down the way right numbers
A: Right.

So community group speakers were increasingly more likely to check the description scheme with the new partner as the experiment progressed. Could this have played a role in the original convergence process by promoting one particular scheme over the others? Closer analysis of the incidence of negotiation in the 19 pre-convergence games suggests that it did not. Firstly, it is interesting to note that the only game where the speakers failed to converge contained negotiation, and secondly, of the ten remaining pre-convergence games that contained negotiation, in only five was there a potential scheme conflict. But perhaps the most striking result was the outcome of these negotiations, since in two of the five conflict cases it had the effect of converting one of the pair from the dominant MATRIX scheme, used in their preceding game, to the non-dominant LINE scheme. So, if anything, negotiation early on in the experiment seemed to interfere with the community convergence process.

This would suggest that the initial checking seen in later games – ex-

[8] The second author first identified all instances and then these were checked by the first author, this produced 98% agreement. Only those cases where both authors agreed are included in the subsequent analyses.

emplified in the extract above – arises from a *prior* assumption that the partner is going to describe positions using some version of the MATRIX scheme but uncertainty about which one they intend to use. This kind of negotiation is not, therefore, necessary for group convergence, even though it does seem to play a role in establishing mutual knowledge of the scheme once the group has converged. The conclusion is reminiscent of Garrod and Anderson's previous claim about the irrelevance of negotiation for establishing stable description schemes between isolated pairs of players.

However, there is another potential explanation for the community's convergence which still has to be considered. It is just possible that players in the community group are responding to the fact that on each new game they encounter a fresh partner and this leads all of them to independently adopt the same rather conservative MATRIX scheme. In effect, this alternative explanation would attribute the community convergence and high degree of inter-speaker co-ordination to a special kind of common salience in choice of the MATRIX scheme.

It was therefore decided to run a second non-community group where each player is exposed to a fresh partner on successive games but the partners are not drawn from the same community. If increasing inter-speaker co-ordination was an indirect response to interacting with fresh partners on each new game, then this non-community group should behave exactly as the community group did in experiment 1. However, if the effect comes from establishing a closed community, as predicted by the convergence account discussed above, then there should be no basis for increased co-ordination.

EXPERIMENT 2: INTERACTING WITH DIFFERENT PARTNERS NOT DRAWN FROM THE SAME COMMUNITY

5. Method

5.1. Subjects and design

Thirty subjects were recruited who were all undergraduates at the University of Glasgow. They were paid at the rate of £1 per game for participating in the experiment.

The subjects were split into five sets, and, for each set, the games were allocated according to arrangement shown in Table 1. So the key player, S1, would play five games overall: first, one with S2, then one with S3, and so forth. Apart from S2, each of the others played one game before encountering S1. So the five key players engaged in five games with a different experienced partner on each occasion, but in a situation where no community could be established. For comparison with the first study, we shall refer to

Table 1
Design of Experiment 2
Control group (2): 30 players of whom 5 play 5 games each with a different partner, who has not played with any of the other 5 lead players. The arrangement for one of the 5 groups is shown

Group 1:	Game level 1	**S1** + S2, S3 + S4, S5 + S6
	Game level 2	**S1** + S3
	Game level 3	**S1** + S5
	Game level 4	**S1** + S4
	Game level 5	**S1** + S6

the games played by the five key players as the non-community group games.

5.2. Procedure

The maze game procedure was exactly the same as that used in experiment 1 and the description exchanges were extracted, transcribed and coded in exactly the same fashion.

6. Results and discussion

In order to compare the speaker's performance in the non-community group with that of the community group in the first study, only the 645 description exchanges which occurred in the 25 games played by the key players were used in the analysis. First, we consider the overall distribution of different description types and then turn to a direct comparison of the degree of inter-speaker co-ordination in the two groups.

6.1. Distribution of exchanges using the different schemes

The overall distribution of different description types across the five key player games is shown in Fig. 6. It is quite apparent from the figure that there is no tendency to converge on the MATRIX scheme. In fact, the most popular type overall is LINE, which was used in 44% of the exchanges.

We can compare the relative proportions of MATRIX descriptions across the five games for the key players in the non-community group with those used by the community group in the earlier study. These were entered into a mixed design ANOVA, with group as a between subjects factor (non-community group Experiment 2 vs. community group Experiment 1), and game (games 1–5) within subjects. It revealed a main effect of group, with $F(1, 13) = 69.84$, $p < .001$, and of game with $F(4, 52) = 4.916$, $p < .01$, which reflect the marked difference in proportions of MATRIX descriptions used by the two groups and the overall increase in proportions across games.

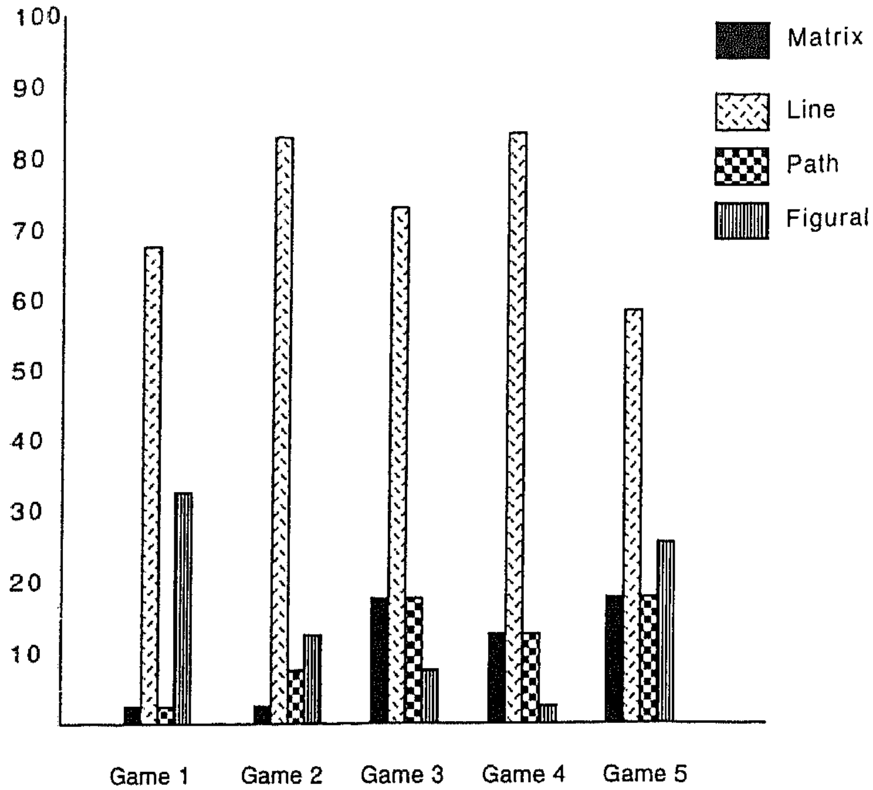

Fig. 6. Average percentage of description exchanges of each type for each player in the non-community group over their first 5 games in Experiment 2.

But it also revealed an interaction between group and game, with $F(4, 52) = 3.85$, $p < .01$, which reflects the marked increase in MATRIX descriptions for the community group ($F(4, 52) = 11.948$, $p < .001$ for the simple effect of game at community) as compared with the non-community players ($F(4, 52) < 1$ for the simple effect of game at non-community).

So it is safe to conclude that there is no evidence for the overall convergence onto the MATRIX scheme, or any other scheme for that matter, among the non-community players. We now turn to the inter-speaker co-ordination measures for the non-community group games.

6.2. Comparison of inter-speaker co-ordination across Experiments 1 and 2

Although there is no evidence for an overall convergence of the non-community group onto one particular scheme, it is also important to establish no increase in the levels of inter-speaker co-ordination. To do this we need to examine the development of co-ordination between the five key players and their different partners, which can then be compared with the related scores for the community group players in Experiment 1.

The sequential probability of matching scores were therefore calculated for all the games played by the key players in exactly the same fashion as in the first study. Figure 7 shows the pattern of results as compared with the community group in Experiment 1.

These data were entered into a mixed design ANOVA, with group (community vs. non-community) as a between subjects factor, and game (1–5) within subjects. This revealed a main effect of group, with $F(1, 13) = 20.77$, $p < .001$, which reflects the dramatic difference in co-ordination

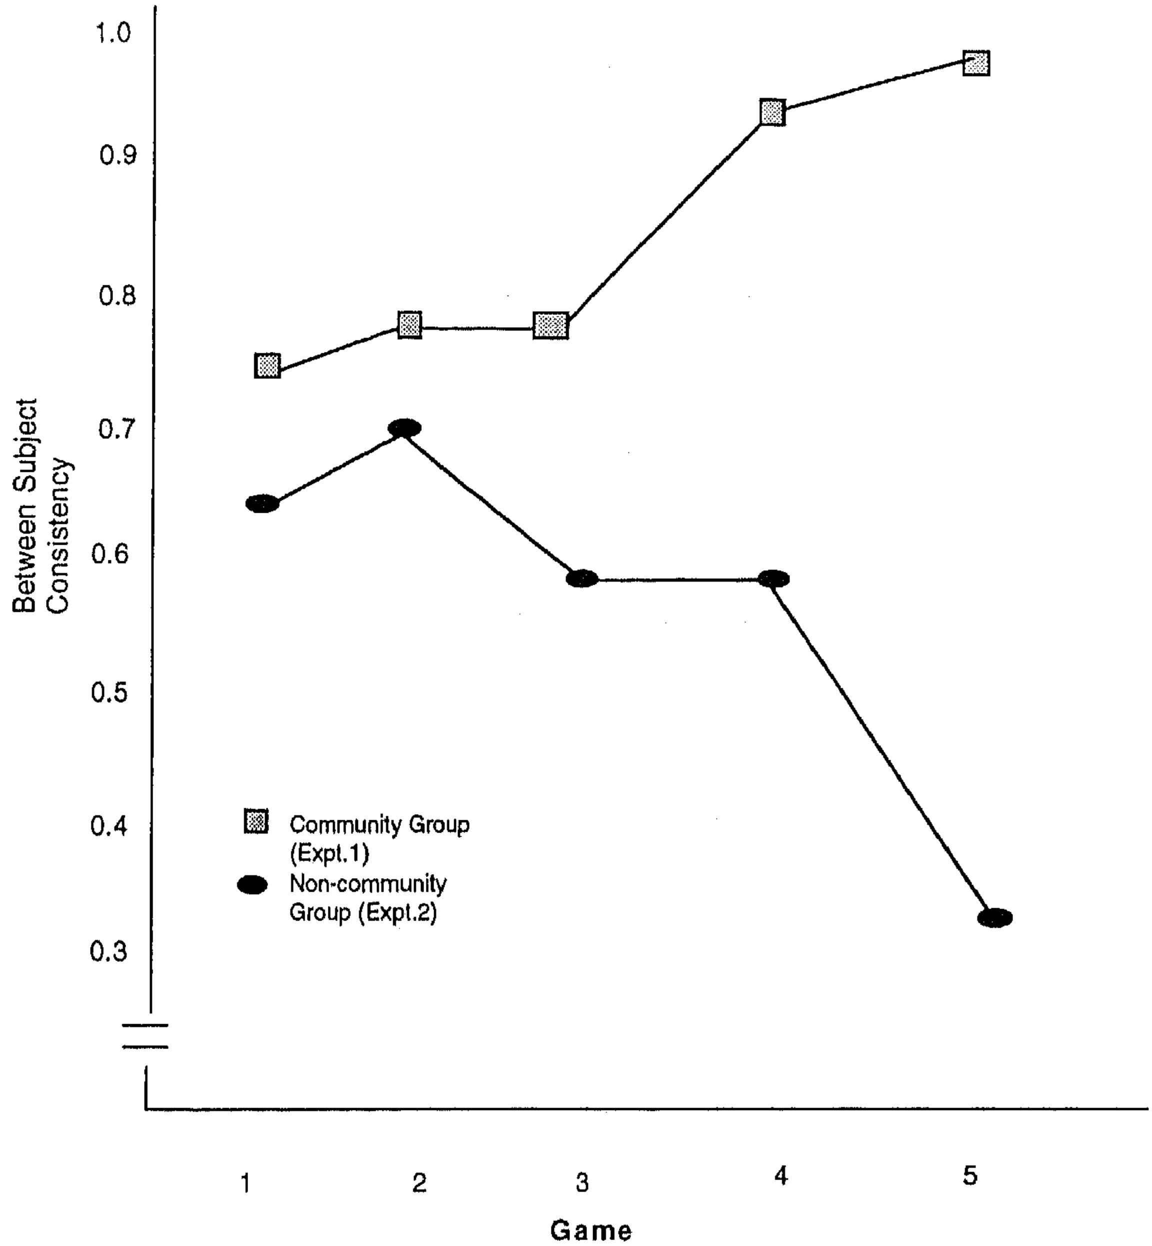

Fig. 7. The mean sequential probability of matching the previous speaker's description in sequences where the speakers change. The consistency measure is shown for games 1–5 of the community group in Experiment 1 and games 1–5 of the non-community group in Experiment 2.

measures between the two (community = 0.847, non-community = 0.581). It also revealed a reliable interaction between group and game, with $F(4, 52) = 3.13$, $p < .05$.

Exploring the interaction further showed that there was a reliable increase in co-ordination only across the first five games for the community group ($F(4, 52) = 3.048$, $p < .05$, for the simple effect), as in the analysis of Experiment 1. However, there was also a reliable difference in co-ordination between the two groups by the time they were playing their fourth and fifth games ($F(1, 65) > 7.22$, $p < .01$ for the simple effects of group at games 4 and 5). So it is clear that speakers in the non-community group are not, in any way, producing the increased measures of inter-speaker co-ordination observed in the first study. In fact, if anything, the trend is towards decreasing co-ordination with greater experience of the game.

This result, together with the overall distribution of different description schemes, clearly rules out an explanation for the group convergence in terms of the general salience of choice of the MATRIX scheme. The convergence process, demonstrated for the community group in experiment 1 must depend upon encountering partners drawn from a closed community.

7. Discussion of Experiments 1 and 2

Taken together the results of Experiments 1 and 2 support a convergence model for community players which reflects the influence of two factors: first, the pressure to co-ordinate the description language between any two speakers, and second, the pressure to retain the most popular scheme overall from their previous encounters. In effect, such a model promotes convergence onto the language scheme which has proved most effective in the past for the group as a whole. The convergence is rapid in a small closed community because each new convert unwittingly infects his or her partner with the scheme and so reinforces its dominance within that community.

This leaves us finally with the problem of explaining how such a group convergence process might eventually lead to the extremely close inter-speaker co-ordination found among the community group members. Of course, as we argued in the introduction, the global co-ordination model predicts elevated levels of co-ordination because convention overrides local precedence and salience constraints, as demonstrated in Experiment 1. But this explanation is not completely satisfactory. In the first place, the high incidence of some form of explicit negotiation in the later games of the community group would suggest that the community $MATRIX_1$ scheme is, at best, proto conventional: even though the scheme is shared, it is not yet recognised as common knowledge in the group. There is also the more basic problem that convention itself is only explained by Lewis at a very abstract level, cast in terms of rational behaviour rather than in terms of psychological mechanism. To give a fuller account we need to consider how the constraints of salience, precedence and convention might be realised in a language processing framework in such a way as to produce the pattern of results demonstrated here.

First we consider how local co-ordination might come about when two context sensitive language processing systems interact, and then, we turn to the problem of explaining the special community group effects observed in Experiment 1.

7.1. *Co-ordination and context sensitivity*

Garrod and Anderson (1987) proposed a general language processing principle which they called output–input co-ordination. It assumes that speaker–listeners are context sensitive to the extent that they draw upon the same local knowledge representation to support both language production and language comprehension processes, and as a result, in conversation each speaker's output will tend to match the last relevant input from their interlocutor. Following such a principle should lead to processing benefits both at the level of the individuals and for the dialogue pair as a whole.

For the individuals there are clear advantages in using the same background knowledge to support both the production and the comprehension process. If, as a listener, you have already interpreted an utterance in a certain way, you will have been forced to take a particular perspective on the topic and to adopt a particular interpretation for the words in the utterance. Hence, when it is your turn to speak, many of the background decisions required in formulating the utterance have, in a sense, already been taken for you. All you have to do is run the generation process over the relevant parts of the decision tree used to interpret recent utterances from your interlocutor (see Garrod, Anderson, & Sanford, 1984, for a discussion of a computer simulation which works along these lines.) This would explain the well established tendency to repeat materials from previous talk (Levelt & Kelter, 1982).

Quite apart from these individual processing benefits, co-ordinating output with input should also be efficient in terms of reducing the overall collaborative effort in dialogue (Clark & Wilkes-Gibbs, 1986). This is because it helps to establish the common ground (Clark & Marshall, 1981). When you interpret an utterance successfully, you can presume that your interpretation matches how the speaker would have interpreted the utterance in that same context. So if you want to facilitate their understanding when next speaking on that topic, your best bet is to match their previous utterance as closely as possible. The co-ordination of comprehension with production processes within the individual conversationalists should therefore promote mutual understanding between them.[9] Furthermore, the whole process does not require that a speaker model his or her partner's knowledge state directly, in the sense of having to hold a distinct model of the other, it simply requires that they each hold a representation of the current state of the discourse as a whole.

Output–input co-ordination would account for the sensitivity to local precedence demonstrated for the isolated pairs in Experiment 1. But how can it account for the occasional influence of salience in the choice of description scheme? Here we must assume that the knowledge representation which underlies internal co-ordination of production with comprehension processes is relatively transient and unstable, possibly being retained in a working memory system subject to interference and decay (see Levelt & Kelter, 1982). With a transient underlying representation precedence alone can only exert a strong influence when there is little intervening dialogue between description exchanges. With more intervening dialogue, other factors, such as salience, might be expected to play an increasingly important role in formulating the description.

This effect of intervening material can be clearly illustrated by considering the dialogues from one of the isolated pairs in Experiment 1. The pair in question adopted a vertical LINE description scheme for most of their exchanges but they never settled down to any one version of the scheme. For instance, the same position could on different occasions be described as "The third row along, fourth down" or "Three rows along fourth down" or "Row three, four down" or "Third row, four down" and so forth. So there was considerable variation in the use of ordinal versus cardinal counting schemes which, as Garrod and Anderson (1987) argued, represent semantically distinct variants of the underlying LINE description scheme.

Taking all the games played by this pair, we were able to identify 146 cases where a description had occurred previously in that dialogue. These were sorted into those which exactly matched the preceding description in terms of the details of the scheme (e.g., use of cardinal or ordinal counting of rows and positions) versus those where some change of sub-scheme had been introduced. They were then organised into two groups: (1) same exchange descriptions, defined as occurring within four dialogue turns of each other; and (2) different exchange descriptions separated by five or more turns. This revealed that only 41% of the different exchange descriptions exactly matched the preceding one, whereas 72% of the same exchange descriptions matched exactly (χ^2 $(1) = 14.87$, $p < .01$). The pattern for this pair therefore indicates that the precedence effect is

[9] For output–input co-ordination to be successful in establishing mutual understanding, it is important that the speakers monitor the success of their utterances. Evidence on description co-ordination in young children's maze game dialogues suggests that 7–8-year-olds co-ordinate their language very closely but without establishing mutual understanding because of failure to properly monitor the communicative success of their descriptions (Garrod & Clark, 1993).

localised, as one would expect if the underlying representation suffered from interference and decay. Among the later community group dialogues there is no evidence of such an effect, since after four games all their exchanges followed the same $MATRIX_1$ sub-scheme.

So it would seem that the kind of inter-speaker co-ordination observed in the isolated pairs may be explained in terms of each individual language processor relying on the same locally stable underlying representation to support comprehension and production. When two such systems interact, they only achieve internal stability when co-ordinated with each other, and, as a consequence, any pair of interacting individuals will always tend to converge in their language use. But how is it possible to account for the extremely high degrees of inter-speaker co-ordination observed in the community group with such a model? We believe that the answer lies in the internal stability and coherence of the underlying representations that speakers can develop in the two cases. Whereas we have suggested that isolated pairs co-ordinate on the basis of transient representations, community members might be expected to build up more stable and coherent representations of the group scheme. One reason for thinking this relates to the differences in each speaker's exposure to the range of different description types and their subtypes.

Establishing a stable underlying representation through the output–input co-ordination process is in many respects like learning a concept from its exemplars. The interacting systems are exposed – either as a result of production or comprehension – to description exemplars, and on the basis of this, they have to derive a common underlying representation sufficiently coherent to support both production and comprehension processes within the overall context of the interaction. It has long been known that the ability to develop stable concepts is very sensitive to the range of variation among exemplars in the learning set (Posner & Keele, 1968). This has also been demonstrated in recent connectionist models of concept learning applied to language acquisition (Plunkett, Sihna, Møller, & Strandsby, 1992).[10]

Community players, unlike isolated pairs, tend to be exposed to a wide range of different descriptions during the pre convergence games. This is a direct consequence of trying to co-ordinate with a number of conversational partners. Whereas, on average each isolated pair speaker produced descriptions representing 1.6 of the four main description types during their first five games, community group speakers produced an average of 2.8 different types of description over the same period, and this difference in range of production is statistically reliable (Mann–Whitney $U = 19$, $p < .05$). Similarly, whereas no isolated pair speaker ever produced more than one sub type of MATRIX description, 6 out of the ten community players produced examples of at least two of the MATRIX subscheme variants at some point

[10] Although the range of variation of exemplar effects is not discussed in this paper, such effects were regularly observed in the simulations (Plunkett, personal communication).

in their first five games. So one reason why community players may establish more stable underlying representations of the group scheme is that they are exposed to a much greater variety of scheme exemplars during the initial pre convergence period.

This increased range of exposure is also true for the non-community group players in Experiment 2, but, since there is no mechanism for group convergence (see above), they have no basis for deriving a stable underlying representation. It is not surprising therefore that the players in the non-community group become increasingly less co-ordinated with their partners as the experiment proceeds, whereas the community group become increasingly more co-ordinated.

Quite apart from this initial difference in range of exposure during the pre convergence games, the community group also encounter their group scheme, once it is established, under a wider variety of interactional contexts than do the isolated pairs. Whereas in the later games, the isolated pairs are only exposed to the established scheme while interacting with one partner, community group players keep on encountering exemplars of the same scheme while interacting with a number of different partners. Again this increased range of exposure in relation to different interactional contexts might be expected to contribute to range effects in stabilising their underlying representation.

So what we are proposing is an account of inter-speaker co-ordination, whether as part of a larger community or in isolated pairs, which treats the process as a kind of dynamic concept learning exercise supported by the output–input co-ordination mechanism. For isolated pairs, who experience a limited range of exemplars, this results in each speaker forming a loose and relatively unstable underlying concept of the joint language scheme. But for speakers who interact as part of a larger community, the learning process involves exposure to a much wider range of exemplars in different interactional contexts. As a consequence, the dominant scheme once established will tend to be more coherent and stable.

This account of co-ordination as dynamic concept learning is also consistent with the group convergence model discussed above. During their initial encounters community group pairs attempt to establish a coherent common concept: an underlying representation which can support co-ordinated production and comprehension within that pair. Not surprisingly, the particular concept on which they converge is determined by what has proved most coherent across the two in their previous encounters with other speakers. When iterated across members of the community this process should eventually lead to all speakers converging on the description prototype or scheme which best captures the range of descriptions reflected across the group as a whole.

Although, some aspects of this psychological model are inevitably speculative, it does offer a unified account of how dyadic conversational interactions among putative members of a community can produce group convergence. Both at the level of the isolated pair and the community as a whole the model is individualistic: convergence comes about because of internal constraints on interacting individuals and when those individuals interact as members of a community those same constraints lead to convergence across the community as a whole.

GENERAL DISCUSSION AND CONCLUSIONS

We are now in a position to address the issue raised at the beginning of this paper: namely, how conversational interaction may influence language change. Our analysis of language convergence during the maze game has identified inter-speaker co-ordination processes that refine the language to suit the local requirements of the communicators. Of special interest is the finding that these co-ordination processes operate not only between isolated conversational pairs but also across a simulated community of speakers. So the experiment provides an example, on a small scale, of a community effecting language change as a consequence of its overlapping conversational interactions. However, before we explore the more general implications of these findings, we need to consider carefully our basis for generalisation.

It might be argued that the community convergence process demonstrated in Experiment 1 may be under the influence of special characteristics of the group itself and could only be clearly demonstrated across a range of different groups. In other words, it is not clear that one can generalise from any one community to another. The reason for thinking this is that community effects may be subject to special interactional factors, such as dominance of one or more individuals within the group, which may violate the sampling assumptions underlying the statistical analyses used in this paper. There are, however, a number of reasons for thinking that this is not a problem in the present case. In the first place, the effect emerged at an early stage in the experiment, long before any one individual or even pair of individuals could have had a sufficiently powerful influence on the inchoate community to produce the degree of convergence observed. As we argued above, the convergence process seems to arise from the mass action of a number of individuals rather than the influence of any particular speaker(s) within the group. In fact, when there was explicit negotiation of the scheme during the pre convergence period – the main way in which speakers might be expected to exert control over the group – it tended to work against group convergence rather than in support of it. So there is nothing to suggest that the group behaviour reflects any property of the whole over and above the sum of its parts.

Another objection concerns generalisation from maze game dialogues to communication at large. Describing positions on a schematic maze clearly strains the resources of conventional language. There are few readily available names for the various locations and so speakers are forced to generate ad hoc descriptions. But we would suggest that it is in just such situations that language change actually comes about. As Glucksberg and Danks (1975) pointed out many years ago language flexibility leads to:

> the development of specific vocabularies within social groups. Lovers, families, schools, trades, professions all develop specialised vocabularies. Photographers speak of *hypo*, psychologists of *shaping*, skiers of *powder*.

We suggest that the kind of language refinement demonstrated in the maze game is of this kind. It reflects the response of the conversationalists to the challenge of communicating efficiently in a locally constrained and novel situation; and these are just the circumstances which stimulate language development in our everyday world of cultural and technological change.

In this broader context, the results from our experiments point to a number of conclusions about the psycholinguistic basis for language change in small closed communities of the kind studied here. Broadly speaking the relationship between control of the language by isolated conversational pairs and control across a community of speakers is in line with the two co-ordination models discussed in the introduction. Isolated pairs co-ordinate their language at a local level on the basis of precedence and salience, whereas speakers in a community come to co-ordinate at a more global level, as if developing a language convention in Lewis's sense.

However, the more detailed analysis of the experiments suggests that the two outcomes of local and global co-ordination arise from the operation of the same basic process of attempting to discover a coherent underlying representation to support both language production and comprehension within each individual speaker. To this extent the community wide co-ordination of the language only represents the initial stages of establishing a convention: producing a sufficiently stable precedent to act as the basis for common knowledge. Taking a more radical view, the experiments suggest that this basic co-ordination process may be an essential precursor to establishing a convention in the first place. In other words, the principle of output-input co-ordination, like Hurford's Saussurean strategy, may be at the root of language development and language change whereas convention, in Lewis's sense, may only be a means of fixing the language once the community has achieved an initial degree of conceptual stability.

Acknowledgments

The experiments reported in this paper were supported by a grant from the British Economic and Social Research Council to the Human Communication Research Centre. The authors gratefully acknowledge Keith Edwards's assistance in creating the Apple Macintosh version of the maze game program used in these experiments. They also acknowledge the helpful suggestions made by three anonymous reviewers of an earlier draft.

References

Anderson, A., & Garrod, S. (1987). The dynamics of referential meaning in spontaneous dialogue: some preliminary studies. In *Communication failure in dialogue and discourse* (pp. 161–183). Amsterdam: Elsevier.

Clark, H.H. (1985). Language and language users. In G. Lindzey & E. Aronson (Eds.), *The handbook of social psychology*, 3rd edn. (pp. 179–231). New York: Harper & Row.

Clark, H.H., & Marshall, C.R. (1981). Definite reference and mutual knowledge. In A.K. Joshi, I.S. Sag, & B.L. Webber (Eds.), *Elements of discourse understanding* (pp. 10–46). Cambridge: Cambridge University Press.

Clark, H.H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, *22*, 1–39.

Garrod, S., & Anderson, A. (1987). Saying what you mean in dialogue: a study in conceptual and semantic co-ordination. *Cognition*, *27*, 181–218.

Garrod, S., Anderson, A., & Sanford, A.J. (1984) Semantic negotiation and the dynamics of conversational meaning. Technical Report, 1. Glasgow University Psychology Department.

Garrod, S., & Clark, A. (1993). The development of dialogue co-ordination skills in schoolchildren. *Language and Cognitive Processes*, *8*, 101–126.

Glucksberg, S. & Danks, J.H. (1975). *Experimental psycholinguistics: an introduction*. Hillsdale, NJ: Erlbaum.

Hurford, J.R. (1989). Biological evolution of the Saussurean sign as a component of the Language Acquisition Device. *Lingua*, *77*, 187–222.

Krauss, R.M., & Weinheimer, S. (1964). Changes in reference phrases as a function of frequency of usage in social interactions: a preliminary study. *Psychonomic Science*, *1*, 113–114.

Levelt, W.J.M., & Kelter, S. (1982). Surface form and memory in question answering. *Cognitive Psychology*, *14*, 78–106.

Lewis, D.K. (1969). *Convention: a philosophical study*. Cambridge, MA: Harvard University Press.

Lewis, D.K. (1975). Language and languages. In K. Gunderson (Ed.), *Language mind and knowledge. Minnesota Studies in the Philosophy of Science* (Vol. 7). Minneapolis: University of Minnesota Press.

Plunkett, K., Sinha, C., Møller, M.F., & Strandsby, O. (1992). Symbol grounding or emergence of symbols? Vocabulary growth in children and a connectionist net. *Connection Science*, *4*, 293–312.

Posner, M.I., & Keele, S.W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353–363.

Schelling, T.C. (1960). *The strategy of conflict*. Cambridge, MA: Harvard University Press.

Schober, M.F., & Clark, H.H. (1989). Understanding by Addresses and Overhearers. *Cognitive Psychology*, *21*, 211–232.