

ISSN 1470-5559

The development of semantic co-ordination in dialogue: the role of direct interaction

Gregory James Mills



RR-07-07

July 2007



**The development of semantic co-ordination in
dialogue: the role of direct interaction**

Gregory James Mills

Submitted for the degree of Doctor of Philosophy

Queen Mary, University of London

2007

Declaration

I declare that the work presented in this thesis is my own work carried out under normal terms of supervision and that the research reported here has been conducted by myself unless otherwise indicated.

Gregory Mills

London, July 17th 2006

The development of semantic co-ordination in dialogue: the role of direct interaction

Gregory James Mills

Abstract

A key problem for theoretical and empirical models of language is to explain how semantic co-ordination in dialogue is achieved and sustained. It is argued that existing theoretical accounts rely on tacit assumptions of a shared semantic code underwriting mutual-intelligibility. These models describe miscommunication as a secondary phenomenon that arises as a complication of trouble-free understanding and do not provide an account of the role of miscommunication in establishing inter-individual coherence in dialogue. By contrast, the empirical accounts provided by the collaborative model and the interactive alignment model both emphasize the importance of interlocutors' opportunities for direct interaction but provide different accounts of how co-ordination is achieved. The collaborative model highlights pair-specific co-ordination processes in establishing mutual belief, while the interactive alignment model emphasizes priming and the matching of representations. To investigate in detail the development of semantic co-ordination and compare the predictions of both models, this thesis describes three 'maze game task' experiments which present participants with the recurrent co-ordination problem of referring to spatial locations. Participants communicate using a novel text-based chat tool which permits fine-grained manipulation of the opportunities for interaction. The first experiment investigates the effects of sequential coherence on semantic co-ordination. The second experiment investigates the role of active participation and the third experiment investigates the role of clarification requests. It is argued that neither model is adequate for addressing the local and global patterns observed across all three experiments. To account for these patterns an alternative model is provided which does not rely on parity of meaning for successful communication and which emphasizes the role of interactive repair as a key process underlying the development of semantic co-ordination.

Submitted for the degree of Doctor of Philosophy

Queen Mary, University of London

2007

Contents

1	Introduction	12
1.1	Overview of this thesis	13
2	Concepts of Communicative Success	17
2.1	Positive characterisations of communicative success	17
2.1.1	The transmission model of communication	17
	The conduit metaphor	18
	Miscommunication and the development of co-ordination	20
2.1.2	Communication as intentional action	22
	Grice's conversational implicature	22
	Searle's speech act theory	24
	Modification of Grice's theory of non-natural meaning	25
	Constitutive and regulative rules	26
	Sincere promising	27
	Indirect speech acts and literal meaning	28
	Network and Background	29
	Summary	29
2.2	Communication in practice: practical criteria of communicative success	32
2.2.1	Austin	32
	Structure of speech acts	33
2.2.2	Wittgenstein	34
	Language games	35
	Multiplicity of language games	36
	Ostension	37
	Meaning as an internal state	39
	Rule-following and judgement	40

	5
2.2.3	Language game failure 43
2.2.4	Ethnomethodology 46
2.3	Conclusions 49
3	Empirical models of (mis)communication 51
3.1	Introduction 51
3.2	Speech act processing 51
3.3	Conversation analysis 54
3.3.1	Repair 58
3.4	Clark and Marshall's model of definite reference repair 62
3.5	Collaborative model 65
3.5.1	Experimental evidence 68
3.5.2	Miscommunication in the collaborative model 70
3.6	Interactive alignment model 72
3.6.1	Miscommunication 75
3.6.2	Emergence of conventions in linguistic communities 76
3.7	Conclusions 78
4	General Methods 81
4.1	Introduction 81
4.2	Initial Pilot Study 82
4.3	Materials 83
4.3.1	Chat tool 83
4.3.2	Chat server 83
4.3.3	Maze game application 84
4.3.4	Procedure 85
4.4	Comparison of media 85
5	The effect of sequential coherence on semantic co-ordination 88
5.1	Introduction 88
5.2	Sequential coherence 88
5.2.1	Experimental hypotheses 90
5.3	Methods 91

5.3.1	Chat tool	91
5.3.2	Participants	91
5.3.3	Procedure	92
5.4	Results	92
5.4.1	Description types	92
	Entrainment	95
5.4.2	Adaptation to the media	96
5.4.3	Effects of reduced sequential coherence	97
5.4.4	Discussion	99
5.5	Conclusions	104
6	The effect of participatory status on co-ordination	105
6.1	Introduction	105
6.2	Participatory status	106
6.3	Methods	107
6.3.1	Chat server	107
6.3.2	Participants	108
6.3.3	Procedure	108
6.4	Results	109
6.4.1	Description types	109
	Transcription results	109
6.5	Discussion	110
6.6	Conclusions	113
7	The role of clarification requests in semantic co-ordination	114
7.1	Introduction	114
7.2	Clarification requests	114
7.3	Methods	116
7.3.1	Chat server	116
7.3.2	Participants	117
7.3.3	Procedure	117
7.3.4	Experimental hypotheses	118

7.4	Results	118
7.4.1	Description types	118
7.4.2	Redundancy and innovation	120
7.4.3	Other measures of CR effects	122
7.5	Discussion	122
8	Discussion	126
8.1	Tacit modification against a backdrop of alignment	128
8.1.1	Granularity of response to CRs	128
8.1.2	Variation within description types	130
8.1.3	The role of variation in interactive repair	133
8.1.4	Summary	140
8.2	Conclusions and further work	142
A	Maze game transcripts	162
A.1	163
A.2	164
A.3	165
A.4	166
A.5	167
A.6	168
A.7	169
A.8	170
A.9	171
A.10	172
A.11	173
A.12	174
A.13	175
A.14	176
A.15	177
A.16	178
A.17	179

List of Figures

4.1	Client chat tool	83
4.2	Example pair of maze configurations	85
5.1	Dual window chat tool	91
5.2	Global distribution of description types in dual and single window dialogue	94
5.3	Distribution of description types in dual and single window dialogue. Further subdivided into the first four games (early) and last four games (late).	95
5.4	Entrainment scores of dual and single window dialogue	97
5.5	Turn formulation time of questions and non-questions in dual and single window dialogue.	98
6.1	Distribution of description types in spontaneous dialogue in the first four games (early) and the last four games (late).	110
6.2	Distribution of description types in response to artificial questions by primary and peripheral participants	111
7.1	Distribution of description types in baseline and clarified dialogue	119
7.2	Distribution of description types of target turns in the first four games (early) and the last four games (late)	119
7.3	Innovation scores of all turns excluding CR responses	120
7.4	Redundancy of responses to ‘Whats’ and ‘Frag’s’	121
8.1	Maze corresponding to excerpt 8.1. Numbers in parentheses correspond to line numbers of excerpt.	130
8.2	Maze corresponding to excerpt 8.3. Numbers in parentheses correspond to line numbers of excerpt.	131
8.3	Maze corresponding to excerpt 8.6. Numbers in parentheses correspond to line numbers of excerpt.	135

8.4 Maze corresponding to excerpt 8.7. Numbers in parentheses correspond to line numbers of excerpt. 138

List of Tables

3.1	Clark and Marshall's 1981 model of definite reference repair	63
4.1	Artificial clarification subdialogue	82
5.1	Length in characters of dual and single window turns	97

Acknowledgements

I would like to express my immense gratitude to my supervisor, Dr. Pat Healey for his unwavering support and encouragement, for introducing me to the study of interaction, especially for introducing me to Wittgenstein's ideas which radically changed my thoughts about language, and also for his endless patience while I held on to unrealistic ideas on how to approach the study of dialogue.

I am also grateful to my second supervisor, Dr. Graham White for his helpful comments and invaluable discussions on the philosophy of language.

I'd also like to thank my friends who endured me stopping mid-conversation and asking them questions about what they meant when they said 'what?' or 'huh?', Rob Bradshaw, Marcus Colman, Arash Eshgi, Jackie Foulkes, Emily Foulkes, Peter Goldsmith, Rita Harrak, Chiranti Kariyawasam and Marc Rollinger. I'd especially like to thank both Alex Barakan and Marc Clement for their humorous encouragement and willingness to be a test-bed for all of my ideas about language and dialogue.

Most of all, I would like to thank my family to whom this thesis is dedicated.

Chapter 1

Introduction

The main question this thesis attempts to address is how co-ordination is achieved in dialogue. Although actual face-to-face dialogue is the primary and most basic site of language use and acquisition, psycholinguistic and semantic models' main objects of inquiry have been predominantly limited to individual texts or sentences in monologic settings. These accounts typically investigate linguistic comprehension or production in isolation, denying the significance of some of the most important aspects of language; in particular that language is adaptable to new contexts of use, that variation is inherent to language and that this necessarily occurs through inter-individual processes in dialogic settings.

The relegation of the importance of actual language use has a rich history which stems partly from the prioritization of written over spoken language due to written language becoming progressively autonomous and persistent; allowing the accumulation and comparison of knowledge, enabling the centralization of bureaucracies and governance, and forging national identities by rendering languages into dialects of the standard written form (Linell, 1982). Associated with this prioritization is the view that the meaning of the words inheres in the text, leading to the notion that words used in communicating in the first instance stand for things or ideas. This conception of language can be traced back via Saussure, Locke and Descartes to Aristotle's nomenclaturism. Ancillary to this conception is that successful communication rests on the identical transmission of meaning:

'To make words serviceable to the end of communication, it is necessary, as has been

said, that they excite in the hearer exactly the same idea they stand for in the mind of the speaker. Without this, men fill one another's heads with noise and sounds but convey not thereby their thoughts, and lay not before one another their ideas, which is the end of discourse and language'. (Locke, 1690)

Central to the claims developed in this thesis is that this conception of language does not readily provide an account of what occurs when interlocutors differ in their interpretations of the meaning of words. This thesis reframes the problem as an empirical question of how problems of interpretation are resolved through inter-individual processes of co-ordination in dialogue, examining how mutual-intelligibility is achieved and sustained and the subsequent effect on the development of co-ordination.

To this end, this thesis draws on philosophical and psycholinguistic approaches to language, but is firmly situated within the language-as-action (Clark 1996) tradition and is primarily data-driven. It draws on experimental paradigms that investigate inter-individual co-ordination in spontaneous task-oriented dialogue as a further development of the experimental program of this tradition. In a sense, the experimental approach adopted by this thesis brings the written-language tradition full-circle; technological advancement has broadened the use of writing from the dissemination of texts to computer-mediated text-based dialogue as a commonplace means of communication, allowing the exploitation of this form of communication as a resource for the experimental investigation of how semantic co-ordination is achieved.

This thesis uses experimental evidence to develop the claim that the structural features of dialogue, including sequential coherence and interlocutors' opportunities for modifying each other's utterances (Healey and Mills, 2006) are of key importance to the development of semantic co-ordination. It provides an account of how the local processes involved in clarification dialogues contribute toward this global development (Mills and Healey, 2006) arguing that miscommunication is one of the progenitors of the development of semantic co-ordination.

1.1 Overview of this thesis

Chapters 2 and 3 present the theoretical background to the experimental investigations carried out in this thesis. Chapter 2 examines theoretical models of communication with the aim of clarifying the relationship between concepts of communicative success and concepts of meaning. Reddy's (1977) account of the 'conduit metaphor' is used to situate the problem of theoretical models'

implicit reliance on interlocutors possessing a shared ‘code’ or set of conventions. This chapter argues that this reliance on shared guarantors of meaning underwriting mutual-intelligibility renders them unsuitable to addressing how co-ordination in dialogue is achieved. This chapter moves from positive characterisations of communicative success of intentionalist models to a consideration of practice, drawing on Wittgenstein’s arguments against the privileged status of reified rules and representations in determining the successful use of language. Wittgenstein’s notion of language games serves to illustrate the inadequacies of variants of the ‘code model’ of communication. This thesis identifies tensions between Wittgenstein’s insights into the inherently dynamic, fluid and revisable nature of language and his assumptions of harmonious practice within a language game, developing the claim that the consideration of practice suggests the importance of examining in closer detail the negative criteria of communicative success. Following Wittgenstein’s recommendations of a detailed descriptive program to establish what the practices are in any given case, the first chapter concludes with the ethnomethodological redefinition of the problem of mutual-intelligibility for both interlocutors and the analyst. The chapter demonstrates how the ethnomethodological program goes beyond Wittgenstein’s artificially constructed language games, providing the empirical and theoretical groundwork for addressing how mutual-intelligibility is achieved.

The following chapter is concerned with empirical approaches to language use that in contrast to the theory-driven models reviewed in the first chapter are resolutely ‘data-driven’. It reviews models that deal explicitly with problems of inter-individual coherence in dialogue, discussing the relative emphasis on positive and negative characterisations of communicative success and the associated concepts of meaning. This chapter focuses on the interactive alignment model of Pickering and Garrod (2004) and the collaborative model of Clark et al. (1996) showing how they directly address the questions raised in the first chapter concerning the dynamic nature of language and how they promise to enhance the analysis of the processes and mechanisms involved in the development of semantic co-ordination. It examines differences between both models’ accounts of the local and global processes underlying co-ordination, describing the experimental evidence used to support both accounts and identifies in the maze game experiments of Garrod et al. the global development of abstract referring expressions used to describe locations in the maze. This chapter diagnoses both models’ lack of a mechanism that is sensitive to semantic differences between kinds of description and argues that these differences in abstraction entail

different levels of co-ordination. This provides an index of semantic co-ordination that contrasts with the collaborative model's focus on iterative cycles of grounding and the interactive alignment model's prioritization of matching of representations.

Having established the theoretical background to the thesis, the following chapters are concerned with the empirical investigation of the development of semantic co-ordination. While the primary co-ordination mechanisms of the interactive alignment and collaborative models emphasize the key importance of direct interaction, evidence for both models is gained indirectly from relatively coarse-grained experimental manipulation of the communicative context. The fourth chapter discusses the approach adopted in this thesis that attempts to address these shortcomings. It describes the modification of the maze game from a spoken task to a text-based computer-mediated task. This requires participants to communicate using a chat tool that permits the fine-grained experimental manipulation of the unfolding dialogue. Selectively interfering with the processes involved in inter-individual coherence allows the direct testing of hypotheses concerning the role of both models' mechanisms in accounting for the global development of abstract descriptions in the maze game.

Chapter 5 investigates the role of sequential coherence in co-ordination. Both models assign importance to the turn-by-turn displays of positive and negative displays of understanding in accounting for how co-ordination is achieved. These displays rely on dialogue coherence provided by turn-adjacency and relevance for their successful deployment. The aims of the experimental manipulation of coherence are twofold. First, to provide a baseline to assess the viability of the combination of the maze game with the chat tool. The second is to observe the local effects of interfering with participants' ability to rely on sequential coherence in order to establish which particular mechanisms are affected and how these local effects are reflected in the global development of semantic co-ordination. The chapter illustrates local and global effects that are not addressed by either model, arguing that the main locus of these effects of reduced sequential coherence lies in participants' opportunities to formulate questions concerning problematic understanding.

Chapter 6 focuses on the effect of participatory status on semantic co-ordination. It reviews existing accounts of co-ordination achieved by interlocutors associated with different levels of active participation and describes the use of the chat tool to manipulate the apparent origin of interlocutors' utterances. The results provide evidence of local and global semantic co-ordination

phenomena that are not captured by the mechanisms of either model. This chapter argues that the results substantiate the conclusions drawn in the preceding chapter, highlighting the importance of participants' opportunities to engage in interactive repair in order to resolve difficulties concerning the use of descriptions.

Chapter 7 investigates the role of interactive repair in semantic co-ordination. While the previous two experiments strongly suggest the importance of signalling and resolving misunderstandings, the evidence is obtained indirectly by hindering participants' opportunities for engaging in repair. The experimental manipulation performed in this task takes the opposite approach of introducing artificial clarification requests into the unfolding dialogue. The chapter discusses the problems concerning the relationship between local and global effects on semantic co-ordination, as the global patterns do not appear to be reflected in the local clarification sequences.

The final chapter draws the preceding chapters together, addressing the conflicting evidence concerning the role of clarification provided by the third experiment. It argues that the unit of analysis adopted to investigate the development of semantic co-ordination is too coarse-grained to capture the local patterns involved in clarification sequences. It provides an account of how clarification requests act as the progenitor of global co-ordination through the systematic exploitation of 'concreteness' which involves participants' fine-grained mutual modification of each other's turns. This thesis examines the possibilities of developing this account using the chat tool and concludes with the implications of these findings for existing approaches to dialogue.

Chapter 2

Concepts of Communicative Success

2.1 Positive characterisations of communicative success

The aim of this chapter is to clarify the relationship between concepts of communicative success and concepts of meaning. Moving from positive characterisations of communicative success through a consideration of practice to negative characterizations of communicative success. It explores the contrasting analyses of meaning provided by information theory, speech act theory and Wittgenstein's language game metaphor, leading to the empirical question of how co-ordination develops through interaction and miscommunication.

2.1.1 The transmission model of communication

Shannon and Weaver's (1949; 1966) information theoretic model of communication is based on the mathematical analysis of electrical signals. Their original model consisted of six elements: an information source that selects and produces a desired message; a transmitter that encodes messages into signals; a channel that conveys the signals; a receiver that decodes signals into messages; a destination where the decoded message arrives; noise that acts on the channel and interferes with the signal conveyed through the channel. Using as an example a telephone conversation, the wire plays the role of the channel, the signals are the modulated currents transmitted through the wire, and noise is the static or interference. In face-to-face conversation the speaker's mouth is the transmitter, the air the channel, and noise is any kind of auditory distraction. Importantly, this code model, when applied to human communication embodies what Reddy (1979)

refers to as the conduit metaphor.

The conduit metaphor

Reddy (1979) proposes that the dominant metaphor used by English language speakers to describe communication is as a 'conduit' in which ideas or meanings move between people. He gives many examples of what at first seem unremarkable everyday accounts of communication going astray, such as (p. 286-289):

Try to get your *thoughts across* better.

None of Mary's *feelings came through to me* with any clarity.

and of remedies to resolve misunderstanding:

You have to *put* each concept *into words* very carefully.

Insert those ideas elsewhere *in the paragraph*.

Reddy argues that these display a common underlying 'conduit metaphor' of thoughts and feelings as physically transmitted between people. Speakers insert their thoughts into words, the words contain the thoughts and act as a conduit conveying them between a speaker and a listener who then extracts the original thought from the words. Reddy proposes that the 'conduit' metaphor has a major and a minor version. In the minor version, thoughts and feelings are treated as being 'ejected by speaking or writing into an external "idea space" where they are reified so that they exist independently... and may or may not "find their way back into the heads of living humans"' (Reddy, 1979, p. 291). Variations of this metaphor are common in everyday folk-linguistic accounts (Haapmaki, 2002; Craig, 1999).

To highlight the assumptions implicit in this metaphor, Reddy constructs an alternative model of communication that he calls the 'toolmakers paradigm'. He supposes there exists a community that lives in a circular compound that is divided into separate sectors. Each sector contains slightly different environments of plants, trees, water etc., and is the living space of an individual. At the centre of the compound is a turntable mechanism that only permits the exchange of special pieces of paper (blueprints) between the inhabitants of the different sectors for the purpose of exchanging instructions with each other. As each sector is enclosed within walls, the only means of communication between members of the community is via this turntable. Under this 'postulate of radical subjectivity' (p. 295), each environment stands for an individual's mental representations

as they are inaccessible to the other members of the community and the blueprints represent the ‘signals of human communication’.

Reddy imagines a situation in which a member sends identical blueprints to the other members of the community which detail how to fashion a tool. The other members attempt to construct the tool from the available materials; some using stone instead of wood, leading to a wide variation in members’ inferences about the applicability of the tool in their respective environments. This leads to many further exchanges using the turntable that challenge each other’s assumptions, resulting in the members ‘raising themselves to a new plateau of inference about each other and each other’s environments’ (p. 295).

Viewed against this backdrop, the conduit metaphor suggests that the toolmakers exchange actual copies of materials from their environment; the contents of messages in the conduit metaphor are to be thought of as reified objects. As a result the conduit metaphor and the toolmakers paradigm differ in their account of what makes communication successful. The former sees communication as predominantly effort-free; understanding can only take place after the message has been transported, and mere exposure to the message should suffice. The conduit metaphor precludes variations in interpretation. With successful communication being the norm, unsuccessful communication becomes a special case requiring an explanation that locates effort in the speaker’s ‘packaging’ of the message.

In the toolmakers paradigm, however, interlocutors’ understanding will necessarily be divergent as they only have partial understanding of each other. Convergence and communicative success can only be reached via ‘continuous effort and by large amounts of verbal interaction’ (p. 295).

Reddy attempts to show how the pervasiveness of the conduit metaphor in the English language has real tangible effects on its speakers, leading to a semantic pathology: the meanings of concepts used to describe the ‘containers’ holding ideas, thoughts etc. are over-extended to include the ideas that they supposedly contain. An example of this is how the word ‘poem’ which is used to stand for both the written text and what he calls its ‘word-sense’. From the toolmakers’ perspective these are two radically different entities, as words do not contain thoughts: a poem in the toolmakers’ compound would only converge after many effortful exchanges.

According to Reddy the perspective afforded by the comparison between these two metaphors leads to a shift away from the the prioritization of the products or artifacts of communication,

which in emphasizing the effortless conception of communication results in the issues presenting themselves as merely technical (storage capacity, speed of access, bandwidth, noise reduction, retrieval etc.) towards an emphasis on the active process of collaborative ‘meaning-making’.

Miscommunication and the development of co-ordination

Although Reddy goes some way towards an account of language that sees meaning as being communicatively co-constructed, it is arguable from a perspective emphasizing the development of co-ordination that his proposed solution is unworkable. The toolmakers paradigm can be seen to be hampered by the same problems as the information theoretic model.

In the toolmakers paradigm, no two conversants have access to each other’s sector, in accordance with the principle of privacy of mental content. The blueprints that are exchanged are supposed to represent the human communication system and are the only means by which the toolmakers can affect each other’s mental content. While it is possible, in keeping with the toolmakers paradigm, for two interlocutors to communicate about their impressions of an object, this account neglects the extent to which they share an environment that affords richer, public ways in which meaning can be negotiated, for example through ostensive definition or through paralinguistic cues. It is significant in this context that Reddy emphasizes the collaborative nature of language and ruminates on the social implications of using the conduit metaphor, when in fact the toolmakers paradigm does not permit any kind of encounter with the ‘other’ except via the exchange of blueprints.

This suggests that Reddy’s description of language is incomplete, as it is not grounded in the complex social lives of its speakers (Wittgenstein, 1958b, *Philosophical Investigations*, henceforth P.I.). For Reddy, the shared practices are nothing more than the shared but solipsistic operation of the communication system, and not the individuals interacting with each other and their environment. It is therefore not evident that the use of a particular metaphor for communication affects understanding to the extent envisaged by Reddy, as ascertaining this would require empirical investigation of the actual practices in which participants made use of the metaphor in instances of misunderstanding. This would also resonate with Sperber and Wilson’s (2002, p. 627) view that ‘in general, folk-linguistic theories about communicative practice have rather limited and peripheral effects on the natural processes of speech and comprehension, where so many of the sub-processes involved are automatic and impenetrable’.

Although Reddy dispenses with the catch-all concept of ‘noise’ to explain miscommunication

and does away with the notion that meaning is ‘contained’ within the messages, his attempt to rid the information theoretic model of objectified meaning is effectively superficial: the toolmakers only engage in a communicative *reconstruction* of meaning, as a speaker’s meaning exists as a reified object prior to the act of communication. Similarly, whilst Reddy’s model does not rely on interlocutors possessing an explicitly shared instruction set or code book, it is supposed that the problem of their possession of a ‘common code’ is resolved by the ‘a priori shared context’ which is different for each interlocutor, yet sufficiently similar to allow communication. However, this model doesn’t specify criteria for sufficiency of similarity and the ultimate guarantee of intersubjectivity is therefore presupposed rather than explained. Moreover, the model invoked by Reddy bears a striking resemblance to Wittgenstein’s (1958a, PI 293) famous ‘beetle in the box’ language game: each person in this language game possesses a box to which only they have access and in which is stored something that they each call a beetle. The term ‘beetle’ also plays a public role in the language game; whenever someone is asked what is in their box, they reply with “‘beetle’”. Wittgenstein argues that in this model of communication the actual contents in the box are irrelevant; each box could have very different contents, yet as no individual has access to another’s box, the contents play no role in the public language game. In effect, the toolmakers’ blueprints cannot be said to have any meaning, as one can ‘divide through’ (Wittgenstein, 1958b, PI 293) by the contents of each environment, thereby cancelling them out. All that could be said about the blueprints is that they function as elaborate instructions for the communication system.

It appears, then, that the toolmakers paradigm, as indeed the information theoretic view of communication being the construction, exchange and reconstruction of messages, relies on interlocutors possessing an objectified shared code, but unlike the ‘code model’ of Jakobson (1960), does not include a ‘contact’ dimension that recognizes the rich, public ways of communicating, or provide an adequate treatment of the role of communicative context. For Reddy, the failings of information theory to provide an accurate account of communication are due to the influences of the conduit metaphor and not due to the inherent reliance on a ‘common code’ as a guarantee of intersubjective understanding.

Further, despite introducing the theme that the development of communication is to a certain extent driven by miscommunication, ‘noise’ is still conceived as extrinsic to the communicative channel, suggesting that in principle it could be overcome.

It is arguable that the toolmakers paradigm of humans communicating within a large com-

pound, sealed off from each other by large walls and exchanging messages at its centre is precisely what Reddy characterizes ‘as replaying the myth of Babel centred around a broadcasting tower’. Ironically, Reddy’s analysis is restricted to the conduit and information-theoretic metaphor, failing to see the multiplicity of metaphors used to describe communication as, *inter alia*, war (Lakoff, 1993), power (Foucault, 1998), resonance (Lewis, 1969) or musical performance (Clark, 1996b) to name but a few.

2.1.2 Communication as intentional action

Implicit in Reddy’s analysis is the theme that the production of messages is an individual, intentional process of transforming thoughts into words, with the recognition that communication does not merely consist in the recovery of a single, literal meaning, but that it depends on context-dependent interpretation of interlocutors’ utterances through a process of inferencing. However, it does not satisfactorily address the question of how this additional recovery of context-dependent meaning is achieved. Although Reddy’s proposed solution brings the active process of ‘meaning-making’ into the foreground, emphasizing the fundamentally social and ‘practical’ aspect of communication, it implicitly relies on the the transfer of meaning between interlocutors. The metaphor does not allow for interlocutors to develop a shared communication system, as the mechanism provided does not allow for interlocutors to converge on a common use for a particular expression, due to their having no opportunities for directly interacting with each other. This prevents the development of a communication system that would allow for the use of words that go beyond the individual’s solipsistic use. Both Grice’s co-operative principle and Searle’s speech act theory emphasize the role of a speaker’s individual intentions in bridging this gap, guaranteeing successful communication.

Grice’s conversational implicature

Grice’s (1975) conversational implicature is an attempt to explain successful communication where there is no regular convention linking an utterance with the intention of the speaker, asking what the underlying rational process is whereby the speaker selects an utterance to convey an implicit meaning (indirect speech act) and also asking how the speaker ensures that the addressee is able to understand this meaning. In Grice’s account, as in Searle’s speech act theory, basic communicative success is presupposed, the question is how do interlocutors go beyond the literal, speaker-independent meaning of an utterance. For Grice, the answer lies in interlocu-

tors knowing and following the same pragmatic rules that are universally accepted as standard behaviour, which he calls the co-operative principle (pp. 45-47):

Make your contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.

Grice gives four maxims, that when observed, fulfill the co-operative principle (henceforth CP)

Quantity: Make your contribution as informative as required.

Do not make your contribution more informative than is required.

Quality: Supermaxim: Try to make your contribution one that is true.

Do not say what you believe to be false.

Do not say that for which you lack adequate evidence.

Relation: Be relevant.

Manner: Supermaxim: Be perspicuous.

Avoid obscurity of expression.

Avoid ambiguity.

Be brief (avoid unnecessary prolixity).

Be orderly.

Significantly, these maxims are not to be thought of as determining interlocutors' actions. Rather, they are 'norms' used in interpretation, which allow the speaker to opt out of following them. In this case the addressee, instead of ignoring the speaker's utterance, may notice the significance of the transgression and use the co-operative principle to infer the speaker's intended meaning. When the speaker's intended meaning goes beyond the literal meaning of the sentence, an implicature is said to be generated.

Grice comments that speakers frequently violate, exploit or even blatantly flout the maxims, listing several types: covert violation (lies), where the speaker intends the addressee to assume that the co-operative principle still holds; overt suspension (jokes, story-telling); overt violations (irony, metaphor) where the hearer doesn't follow the maxim of truthfulness, but is still assumed to be following the supermaxim of quality. So, when faced by apparent violations, the hearer assumes that the speaker is still obeying Grice's CP at a deeper level, enabling the hearer to comprehend the speaker's intention.

Grice's CP is used by Levinson (1983) to argue against conceiving of communication as being reducible to a set of conventions, since the CP shows that whenever there is an expectation

or conventional way of expressing a message, participants are able to exploit this expectation in formulating messages whose 'meaning' goes beyond what is conventionally meant by them.

Searle's speech act theory

Searle, who was Austin's (see below) student in the 50's, performed a formalist revision of Austin's work, starting with the premise of individual intentionality underlying all instances of communication. Searle's enquiry is an attempt to discover the 'necessary and sufficient conditions' (Searle, 1975) for the successful performance of a speech act; semantic rules that are constitutive of any communicative action and hence determine the success of their use. While for Austin, the securing of 'uptake' is essential, Searle's unit of analysis is '...not the symbol or word or sentence or even the token, but rather the production of the token in the performance of the speech act that constitutes the basic unit of linguistic communication' (Searle, 1972, pp. 136-137). Thus successful performance of a speech act only requires a single speaker.

Searle's speech act consists of four separate elements, in the attempt to distinguish between an utterance and its meaning, divided as follows:

Utterance act: The actual production of sounds.

Propositional act: Consists of referring acts and predicating acts that deal with the subject matter and the propositional content.

Illocutionary act: e.g. promising, questioning, ordering. A function-indicating device dictates how the propositional content is to be taken through word order, stress, intonation, etc. Its recognition by the addressee is deemed the illocutionary effect. The function indicating device is not always functionally necessary, as the communicative context often provides this information.

Perlocutionary act: The attempt by the speaker to bring about certain actions by the addressee (perlocutionary effect).

Searle's utterance act has to be accompanied with a propositional act for it to have meaning. Although purely referential speech acts are permitted in this model, this reduction of meaning to the transfer of propositional information is essentially a re-alignment with the transmission model of communication. Searle's account is an attempt to specify how the idealized formal objects constituting speech acts dictate their successful use.

Modification of Grice's theory of non-natural meaning

To this end, Searle draws on Grice's (1957) theory of natural and non-natural meaning that seeks to explain the meaning of sentences and utterances by reference to a speaker's intentions distinguishing between two types of meaning: natural and non-natural (pp. 378-379):

Natural meaning: Black anvil clouds meaning that a thunderstorm is close.

Non-natural meaning: Three rings of the bell mean the bus is full.

Grice points out that it makes no sense to say sentences of the form 'Those black anvil clouds meant that a thunderstorm was close, but there was no thunderstorm', whereas it would be permissible to say 'The bell rang three times to indicate that the bus was full, but in actual fact it wasn't'. The link between the ringing of the bell and the bus being full is a non-natural one, and exists only due to a convention. Grice is concerned with the non-natural sense of meaning, and performs a reductive analysis to derive an explanation of non-natural sentence-meaning in terms of the speaker's intentions. In doing so, he refines the notion of communication as intentional action by positing that a communicative intention is a reflexive act: the intentional effect of the speaker is achieved in the addressee by means of the addressee recognizing that the latter is intended to recognize it: 'For some audience A, U intended his utterance of x to produce in A some effect (response) E by means of A's recognition of that intention' (Grice, 1969, p. 160). Having accounted for speaker's meaning, Grice states that the meaning of a sentence can only be said to have the same meaning as the speaker's meaning if there exists a convention that links it with the speaker's utterance.

Searle's objection is that this does not differentiate between illocutionary and perlocutionary effects. An illocutionary effect is part of the illocutionary intention that is fulfilled merely by the audience understanding what was said. Secondary effects, such as amusing, embarrassing or impressing the audience belong to the perlocutionary effect.

Searle (1972, pp. 229-230) uses the example of an American soldier who is captured by Italian soldiers and attempts to convince them that he is German by saying 'Kennst du das Land wo die Zitronen blühen?', he argues that, unlike Grice who would claim that the sentence means something like 'I am German', the sentence still has a conventional meaning of 'Do you know the country where the lemons blossom?'. Searle states that 'Meaning is more than a matter of intention, it is also a matter of convention' and in order to incorporate the notion of conventional

meaning adds to Grice's account the conditions that the speaker's intentions are illocutionary and that the speaker intends to respect the conventions that govern how the expression is used in the relevant language. Searle's account of utterance meaning is equivalent to Grice's non-natural meaning with an illocutionary effect as opposed to a perlocutionary one. The result is a divorce of sentence or literal meaning and speaker's or utterance meaning, the separation that he relies upon to explain indirect speech acts (see below).

Another consequence of this divorce is that the applicability of a speech act does not rest on previous contexts of use (as posited by Austin). Instead, the context of use will clarify the intention underlying the speech act: saying something akin to 'I will be there tomorrow' in a situation where the audience is in need of help is an act of promising, there being no need to utter 'I promise that I'll be there tomorrow'. The context makes it clear that the utterance is a promise (Lewis, 1969).

However, this modification renders the performance of a speech act even more solipsistic, since Grice's theory of non-natural meaning allows for the possibility of communication without there existing a convention for expressing the message intended by the speaker (Levinson 1983), but this is explicitly disallowed by Searle.

So, even though Searle's speech act does not require uptake, a speaker's intention incorporates a representation of the audience's understanding which is constituted by the formal rules that are the conventions of the language shared by the speaker and the audience. In addition to pre-supposing a shared level of linguistic meaning, it relies on both interlocutors sharing the same knowledge about the rules that govern the illocutionary acts of their language. This shared knowledge of literal meaning, rules and contexts of use amounts to the 'common code' of the archetypal transmission model, as it is a pre-requisite for a successful communicative event.

Constitutive and regulative rules

Searle distinguishes between constitutive and regulative rules that he sees underlying every speech act. The former define the nature of a speech act, and the latter regulate their use. Most of Searle's analysis is conducted in terms of four major structural conditions. The first is the propositional content condition (the combination of referring and predicating). For example, some speech acts, such as apologies, require the propositional content to be formulated in the past tense, others in the future. The second kind are preparatory conditions and are those that need to hold before a speech act can be made properly, e.g. for an apology to be made, some

event has to have occurred that was considered harmful by the addressee. The third is the sincerity condition that is satisfied when the speaker believes their own claims. Searle does not consider this necessary for successful completion of a speech act. The fourth condition is what Searle calls the essential condition that occurs when the speaker intends the utterance to be recognized by the listener as an instance of a conventional speech act and also intends to be placed under the obligations that the speech act entails (Searle, 1969).

Sincere promising

Using the four conditions described above, Searle embarks on an analysis of ‘sincere promising’, yielding the following constitutive rules:

1. Normal input and output conditions obtain.
2. The speaker expresses the promise in the utterance.
3. In expressing the promise, the speaker refers to a future promise act to be performed by the speaker.
4. The future act must be of benefit to the hearer.
5. It isn’t obvious to both speaker and hearer that the promised act is going to be carried out by the speaker irrespective of the promise (Searle introduces a principle of least effort: maximum illocutionary ends with minimum phonetic effort—see section 3.5.1 for the principle of least *collaborative* effort).
6. The speaker intends to carry out the promised act.
7. The speaker intends that his utterance will place him under an obligation to carry out the promised act.
8. The speaker intends that the utterance will result in the listener believing that conditions (6) and (7) obtain, due to the recognition of the intention of the speaker to produce that belief. This recognition is to be achieved by means of the recognition of the utterance as one conventionally used to produce such beliefs.
9. The semantic rules of the dialect spoken by both are such that the promise is correctly and sincerely uttered if and only if conditions (1)-(8) obtain.

Rule (8), which Searle dubs the ‘essential’ condition, is the modification of the Gricean account that emphasizes the importance of an existing ‘convention’ that enables the successful performance of a speech act.

From these nine conditions, he extracts rules for the use of the function indicating device (P) : (1), (8), (9) apply to all kinds of illocutionary acts, the rest only apply to acts of promising.

Rule 1: P is to be uttered only in the context of a larger stretch of discourse which refers to a future act by the speaker.

Rule 2: P is to be uttered only if the hearer would prefer the speaker carrying out the act to not carrying it out.

Rule 3: P is to be uttered only if it is not obvious to both that the speaker would carry out the act in the normal course of events.

Rule 4: P is to be uttered only if the speaker intends to carry out the future activities.

Rule 5: The utterance of P counts as the undertaking of an obligation to do A (essential, constitutive).

Indirect speech acts and literal meaning

In attempting a classification of all speech acts, Searle (1979, pp. 30-57) describes twelve ways in which speech acts can differ from each other, introducing inter-alia the notion of ‘direction of fit’ that can be either word-to-world or world-to-word. He draws on this taxonomy in investigating indirect speech acts. One of the examples he gives is someone saying ‘You’re standing on my foot’. Taken literally, this person is uttering an ‘assertive’ speech act whilst indirectly issuing a ‘directive’. Searle remarks that indirect speech acts fit a variety of moulds of one speech act type couching another, claiming that one of the main reasons they are employed is due to politeness. Literal meaning and speaker meaning come apart in indirect speech acts: an utterance’s literal meaning such as ‘I would like some chocolate’ is an assertive about what the speaker wants, whereas the speaker’s meaning is a ‘directive’, ‘give me some chocolate’. For Searle, literal speech is the more natural, as the sentence says ‘nothing more and nothing less than what the speaker means it to say’.

His detailed examination and account of indirect speech acts is complex, introducing the notions of ‘Network’ and ‘Background’ (see below) as well as drawing on Grice’s co-operative principle, and will not be reproduced here.

Network and Background

In addition to conventions, Searle's later 'Network' and 'Background' introduce even more entities that are purportedly shared by interlocutors (Searle, 1983). The Network is a mesh of interconnected intentions that any speaker's intention must refer to: for someone to announce their intentions to run for President, the candidate would 'normally believe that the country has elections [and] desire that he receive the nomination of the party etc.' (p. 140). There are also subsidiary intentions as well as 'hopes and fears, anxieties and anticipations, feelings of frustration and satisfaction'. Following the strands of these intentional states supposedly leads via unconscious beliefs such as 'elections are held at or near the surface of the earth' (Searle, 1983, p. 142) to a 'bedrock' (cf. Wittgenstein, 1958) of mental capacities that Searle characterizes as pre-intentional states that constitute the Background. Searle states that the boundary between Network and Background is inherently vague and further divides the latter into a deep (human biological makeup) and local Background (local cultural practices). Therefore, a sentence such as 'The cat is on the mat' can only be evaluated against a Background knowledge of the properties of solid objects, gravity, etc.

With these two related concepts, Searle is positing that in addition to intentional states and their associated propositional content there are also non-intentional practical 'know-how' skills, as well as some overlap between the two. Searle's idealist conception of speech acts thus ultimately relies on a notion of 'bedrock' practices and a propositional view of language. Although Searle's Background is not a set of propositional representations, when the Background manifests itself, it does so by giving rise to intentional states that necessarily have propositional content.

Summary

In abstracting away from use, Searle removes the natural variability of real world speech and unlike Austin, who was uneasy with the distinction between perlocutionary and illocutionary acts, relegates the perlocutionary act to a peripheral role.

Searle's and Grice's concepts of meaning both rest on the foundations that language primarily serves the purpose of representing and communicating about thoughts and states of the world. In both accounts the relationship between a speaker's intention and the subsequent reaction by the hearer is seen as simple and causal: communication occurs via a conduit of mutually shared idealized conventions; perlocutionary effects as well as conversational implicatures can only occur when the hearer first understands the literal meaning of the utterance. This ignores any

effects that the credibility, role or status of the speaker might have on the hearer, as well the role played by rhetoric. As a result, it divorces interlocutors' interactional goals from the illocutionary act, e.g. affecting interlocutors' self-image or maintaining social bonds.

By modifying Austin's speech act theory to include a propositional act, Searle's speech act becomes a representational model of communication that contains a representation of the audience's projected understanding that is constituted by the formal linguistic and pragmatic rules of the language that a speaker and listener are both presumed to share prior to the act of communicating. In abstracting away from actual language use and positing these idealized conventions, he removes the natural variability of real world speech; Searle's enquiry breaks away from its Austinian origins to become a study of these idealized conventions rather than an empirical concern with the characteristics of ordinary language use. Although speech act theory is primarily concerned with what a speaker intends with an utterance in performing it ('parole', Saussure, 1974), as opposed to what a given speech act means in a particular language ('langue'), Searle firmly situates his notion of speech acts as a systematization of the latter, stating that 'an adequate study of speech acts is a study of langue' (Searle, 1969, p. 17). Consonant with this is Searle's exclusion of perlocutionary effects, which belong to the domain of the former.

While the underlying guarantee of mutual-intelligibility is provided by 'the semantics of a language [being] regarded as a series of systems of constitutive rules and that illocutionary acts are acts performed in accordance with these sets of these rules' (Searle, 1972, p. 380), the conventional rules are static and pre-given. This precludes any analysis of how, on communicating with interlocutors who have different conceptions of a speech act, interlocutors might revise their own understanding of the rules that constitute and regulate their use. Thus, it gives no account of how these illocutionary acts might emerge out of use.

Similarly, when using Grice's normative CP, both speaker and addressee need to be sensitive to each other's backgrounds, competences, preferences, etc., as people differ in what they perceive to be relevant or truthful according to individual and cultural differences (Sperber and Wilson, 1986, 2002), for example a politics professor explaining the presidential election to a class as opposed to a child. Grice's maxims do not allow conversants to mutually define what they consider truthful or perspicuous and don't take into account the possibility of interlocutors of a higher status enforcing their interpretation of the maxims (Sarangi and Slembrouck, 1992). Also, utterances that seem to flout one of the maxims of the CP often lead to the initiation of

repair instead of leading to a search for non-literal meaning.

Searle's speech act is even more individualistic. All that is required is that the speaker has a representation of the hearer's understanding. In addition to pre-supposing mutual-intelligibility, Searle's account does not take account of collaboration in communication:

It is not obvious to both [a speaker] S and [a hearer] H that S will do A in the normal course of events (Searle, 1969, p. 59).

However, it is relatively easy to find counter-examples. It may be quite clear to H that S will do A, but H forces a promise to underscore the importance H has attached to A. Similar counterexamples can be found for most of Searle's conditions, given the exclusion of 'marginal, fringe and partially defective [promises]' (p. 55) ; this begs the question 'How does it become obvious to both S and H?'

A promise is defective if the thing promised is something the promisee does not want done; and it is further defective if the promisor does not believe the promisee wants it done(Searle, 1969, p. 58).

Searle treats criteria such as 'good' and 'harm' as if they were robust intersubjective concepts that are readily accessible to both participants, ignoring that what counts as promising might differ depending on the context of utterance, and furthermore ignoring the protracted collaborative preparatory negotiations that are often necessary to arrive at a mutually satisfactory definition of promising before the speaker utters 'I promise...' A solution provided by Clark (1983) is to replace the speech act with a 'dialogue act'.

Although Searle recognizes that speech acts typically occur within a sequence of other interleaved speech acts from multiple participants (i.e. in conversation) that are situated within a social setting: 'P is to be uttered only in the context of a sentence (or larger stretch of discourse)' (Searle, 1969, p. 60), and also in his taxonomic account of 'differences in relation to the rest of the discourse' (Searle, 1979, p. 6), his account is restricted to an atomistic analysis of speech acts in monologic settings rather than actual instances of utterances produced sequentially by speakers in a joint setting (Schegloff, 1988, 1992). A further consequence of the use of Searle's formalization in describing communicative sequences is that it leads to a single 'determinate' structure of each communicative act. The preparatory conditions are logically prior to the act of

communicating. For example, it is presumed that one of the preparatory conditions for felicitously giving an order is that the speaker should be in a position of authority over the addressee. This does not allow for a situation in which a speaker establishes their authority in giving an order, which in addition to treating 'authority' as an outcome of the communicative encounter would necessitate the inclusion of 'uptake', as the outcome depends on whether the response is compliant or challenging. Superficially the speech act can be performed in a variety of contexts, yet the core semantic rules for their application remain invariant, precluding any changes that might occur within a speaker, or more importantly within language as a whole.

2.2 Communication in practice: practical criteria of communicative success

2.2.1 Austin

Austin's theory (1962) emerged out of his criticism of the positivist account of language that underlies the transmission model of communication: meanings are contained in words and utterances that serve the purpose of representing 'states of affairs' in the world or mental content. On this 'constative' view, only statements whose truth-conditional value can be evaluated are seen as meaningful. Austin rejects this account, indicating that language is not only used to represent 'states of affairs' but also to effect a change in the world, that is to perform an action within a social context such as naming a ship or betting: 'I hereby...'. He makes the distinction between constative and performative uses of language in terms of their 'direction of fit' (Austin, 1953, p. 234), remarking that constatives fit 'words to the world', whereas performatives 'fit the world to one's words' and do not have truth-values. Instead, performatives are either felicitous or infelicitous (mutually exclusive), depending on whether they bring about this change. For some performatives to be felicitous, they have to be uttered within an appropriate context that is established through convention: to promise, it suffices to utter the words 'I promise...', whereas marriage vows only count as such in the presence of a priest or registrar.

The factors that determine the felicity of a performative are both 'internal' and 'external' to the interlocutor: in the former case, a performative is abused when uttered without sincerity or is not fulfilled, and in the latter the performative is said to 'misfire' when the act fails to meet the criteria required by the convention or is not fully carried out.

Austin notes that the different kinds of misfire and abuse often overlap and are also 'subject to certain whole dimensions of unsatisfactoriness to which all actions are subject but which are

distinct or distinguishable from what we have chosen as infelicities', by which he means situations in which the utterance was coerced or made mistakenly. For Austin, social conventions are inherently vague as there will always be marginal cases which cannot be determined from their previous application, giving the example of baptising a dog. Unlike Searle, Austin does not claim that there are static fixed rules (conventions) governing the applicability of speech acts; his solution to the problem of 'open texture' (Waismann, 1951) is empirical and similar to Wittgenstein's in that he sees the contexts surrounding previous performances acting as the background against which the applicability of speech acts is determined. However, Austin's account does not emphasize the importance of these marginal cases in the development of new conventions.

After a comparison of the logic between constatives and performatives he concludes that most, if not all utterances have both a performative and a constative aspect. Instead of rejecting the dichotomy, he synthesizes constitutive and performative acts to form a 'speech act'.

Structure of speech acts

Austin divides the speech act into three separate units of analysis: locutionary, a meaningful utterance; illocutionary, a meaningful utterance with a *conventional* force; perlocutionary, a meaningful utterance with a conventional force bringing about a non-conventional effect. The locutionary act is further divided into the phonetic, phatic and rhetic acts which are, respectively, the production of noises, the production of sentences and the production of meaningful speech to bring about a certain effect. The same theme can be uttered on different occasions, yet be associated with different rhemes. Likewise, different rhemes, when used in similar contexts and according to similar conventions can be thought of as 'rhetically equivalent' acts. This division of the locutionary act moves away from an idealized view of language, as unlike Searle, it doesn't posit that the utterances have any propositional content that is prior or essential to the other possible effects. For Austin, 'The total speech act in the total speech situation is the only actual phenomenon which, in the last resort, we are engaged in' (Austin, 1962, p. 52). This can also be seen in his not adopting a transmission model, as in place of saying that information or meaning is *conveyed*, he prefers to see it as an *effect*.

Locutionary and illocutionary acts are intrinsically bound, for performing an illocutionary act is necessarily performing a locutionary act, and in performing a locutionary act, the speaker is always asking a question, or confirming or warning etc. In much the same way that phatic and rhetic acts are not the consequence of phonetic acts, illocutionary acts are not to be thought

of as the consequence of locutionary acts: saying 'I promise...' does not causally bring about the act of promising, rather it is *constitutive* of the convention of promising; the person doing the promising is performing the act *in* saying something as opposed to *of* or *by* something. This moves the performance of a speech act away from the Searlean view that the speaker actively 'participates in the constitution of the act as an illocution' (Cobley, 1996, p. 254).

For an illocutionary act to be successfully performed, the recipient has to understand the meaning and strength of the illocutionary force, i.e. 'uptake' must be secured. This emphasis on the relational character of speech acts is not made by Searle who treats communicating as an individualistic enterprise.

Austin concedes that in some cases it might not be possible to determine whether an act is perlocutionary or illocutionary, giving an example of a man swinging a stick. In one sense, this could be construed as an illocutionary act of warning or in another sense as an accidental side-effect of violent intent where the warning would be perlocutionary. In this case the decisive factor is whether stick-swinging is conventional for both parties, to which Austin (1962, p. 119) remarks 'it is difficult to say where conventions begin and end'. So for Austin, conventions are not clear-cut idealized entities, and although he does attempt a taxonomy of different illocutionary acts, his approach attempts to avoid idealization and endowment of formal logical properties, admitting that 'there are still wide possibilities of marginal or awkward cases, or of overlaps', as well as admitting that much further work is required.

2.2.2 Wittgenstein

In his later work, Wittgenstein, like Austin and Reddy, came to believe that 'positivistic' 'truth-functional' accounts of language were fundamentally flawed. He also sought to correct this through a detailed consideration of language in ordinary use. Wittgenstein's consideration of the actual practical use of language provides a radically different conception of the positive criteria that determine the (in)correct use of an expression. Wittgenstein's emphasis on the fundamentally social nature of language, highlighting how language is inextricably interwoven within everyday practical activities, runs counter to the search for idealized and systematic rules that interlocutors draw on 'wilfully and intellectually' (Saussure, 1974, p. 14). From the idealized vantage point, the 'confused mass' of Saussure's 'parole' or Searle's 'marginal cases' of language use (Searle, 1969) cannot be put into any 'category of human facts' (Saussure, 1974, p. 14). This fosters a reification of the static elements of communication and their separate study in isolation with the

aim of achieving objective knowledge that is independent from their context of use. When these objects of knowledge are then used to explain real-world phenomena, they give rise to Jakobson's code, Searle's conventions, and Saussure's 'langue', all of which are entities that interlocutors are pre-supposed to share in order to allow communication.

Of most relevance to the themes introduced in this thesis are Wittgenstein's criticism of theories of language that ground the correctness of use in objects, whether external or internal, his demonstration of how correct application cannot be taught ostensively and his arguments against seeing communication as behaviour governed by static reified rules.

Language games

In order to draw attention from thinking of language consisting of a system of meaningful signs to how language is embedded within the active social life of its speakers, Wittgenstein introduces the notion of 'language games'. A language game is not to be thought of as some definite atomic entity, such as a speech act, rather, it is a metaphor that Wittgenstein uses to highlight the practical use and learning of language in simple concrete cases.

'I shall in future again and again draw your attention to what I shall call language games. These are ways of using signs simpler than those in which we use the signs of our highly complicated everyday language. Language games are the forms of language with which a child begins to make use of words. The study of language games is the study of primitive forms of language or primitive languages. If we want to study the problems of truth and falsehood, of the agreement and disagreement of propositions with reality, of the nature of assertion, assumption and question, we shall with great advantage look at primitive forms of language in which these forms of thinking appear without the confusing background of highly complicated processes of thought. When we look at such simple forms of language, the mental mist which seems to enshroud our ordinary use of language disappears. We see activities, reactions, which are clear-cut and transparent. On the other hand we recognize in these simple processes forms of language not separated by a break from our more complicated ones. We see that we can build up the complicated forms from the primitive ones by gradually adding new forms' (Wittgenstein, 1958a, p. 17)

Wittgenstein uses language games as a unit of analysis to elucidate the notion that language is not simply to be thought of as a system of signs that can be considered in isolation and ab-

stracted from their actual use. In contrast, language games are used to demonstrate the ‘in situ’ use of language, providing simple but internally consistent examples of patterns of structured social activity in which the use of language in carrying out that activity is embedded, allowing Wittgenstein to provide a detailed descriptive account. Importantly, through this avoidance of abstracting from the communicative context, the use of language games cuts across the distinction between illocutionary and perlocutionary effects of language:

‘We are so much accustomed to communication through speech, in conversation, that it looks to us as if the whole point of communication lay in this: someone else grasps the sense of my words which is something mental: he as it were takes it into his own mind. If he then does something further with it as well, that is no part of the immediate purpose of language’ (Wittgenstein, 1958b, PI 363)

Observing language use in situ, for Wittgenstein, is in the first instance describing the activities in which it is embedded, necessarily including the analysis of perlocutionary effects. Indeed, from a language game perspective, this distinction between illocutionary and perlocutionary acts appears superficial, since it cannot be excluded that what is an illocution in one culture is not a perlocution in another (Duranti, 1997).

Multiplicity of language games

Wittgenstein criticizes the Fregean and early Wittgensteinian view that there is an enumerable set of basic kinds of sentences, e.g. questions, assertions and commands. This view was also held by Searle and to a lesser extent by Austin, whose approach to systematization eschewed the use of logical properties to perform this classification, instead thinking he could arrive at a taxonomy from inspecting all the performative verbs:

‘We should not despair too easily and talk, as people are apt to do about the infinite uses of language. Philosophers will do this when they have listed as many, let us say, as seventeen; but even if there were something like ten thousand uses of language, surely we could list them all in time. This, after all, is no larger than the number of species of beetle that entomologists have taken the pains to list’ (Austin, 1970, p. 234)

For Wittgenstein, there are ‘countless different kinds of use of what we call “symbols”, “words”, “sentences”’. These uses are not fixed, ‘new types of language, new language games...

come into existence, and others become obsolete and get forgotten... bring[ing] into prominence that speaking a language is part of an activity or of a life-form' (PI 23). In much the same way that language games are used to draw attention to the way in which language is embedded in non-linguistic practices, the notion of language as a 'life-form' is used to show how language is embedded within and shapes the cultural practices of communities; it is impossible to understand language by abstracting it from its natural context into some logical domain. The mastering of language involves a never-ending process of learning, practicing and becoming acculturated to new language games of a community. The similarities between the language games that constitute a language can be thought of as 'family resemblances' (PI 67). Wittgenstein uses the example of games; there is not one particular trait common to all games, yet they are grouped together in much the same way that a rope consists of individual strands of fibre, its strength relying not on a single fibre running through the whole thread, but on the 'continuous overlapping of these fibres' (PI 67).

Wittgenstein thus rejects the attempt to find a formal and universal set of rules that govern communication. He sees speakers' linguistic competences as emerging gradually in a never-ending process, their skills developing as they practice language games and encounter new ones. In fact, it would not make sense for Wittgenstein to speak of two language speakers sharing the same universal competences, as they would necessarily have very different interaction histories.

This contrasts for example with Searle's detailed analysis of promising that sees promising as a single entity and focuses on specifying the internal, intentional, logical conditions and rules. Differences between Wittgenstein's and Austin's approaches to investigating the actual uses of language are also apparent: Austin was concerned with capturing the *total* speech act, an enterprise that from a Wittgensteinian perspective is unachievable. While the use of the language game brings the public and observable practices of interlocutors into the foreground, capturing the total speech act would require elucidating the complete interaction histories of all the interlocutors involved. In the constructed examples of simple and minimal language games, the communicative behaviour is 'open to view'.

Ostension

Implicitly in Reddy's account, as well as in Searle's and Grice's, is the assumption that underlying most uses of language is a propositional content that serves to represent 'states of affairs' of the world. Thus, in Searle's treatment, the following speech acts: 'Could you please put the book

on the table?', 'The book is on the table', 'The book was on the table before someone took it' all share the common semantic content of a book being on the table. A common assumption underlying such 'picture' theories of language is that correct use of terms such as 'book' can be taught through ostensive definition. Wittgenstein presents a language game, an imaginary concrete instance of everyday language use: a shopkeeper takes an order via a series of simple comparisons between a shopping list and labels on a set of drawers and then takes the right number of items out of the drawer by following a mechanical process of reciting the cardinal numbers. Wittgenstein asks the question of what the meaning of 'five' is, saying that 'No such thing was in question here, only how the word "five" is used' (PI 1) showing that it does not make sense to ask for which object the word 'five' stands. What is significant, is what the shopkeeper actually does; highlighting how language is inextricably interwoven with everyday activities. To further examine the assumptions underlying the view that the criteria for successful use can be taught through ostension, he gives an example of a 'complete and primitive' language that is used by a builder and an assistant for the purpose of fetching blocks, pillars, slabs and beams. In imagining how this simple language might come to be learnt, he shows the ostensive definition of learning is insufficient, as it presupposes that the learner already comprehends the notion of naming and does not account for the learning of the process of fetching stones. 'One has already to know (or be able to do) something in order to be capable of asking a thing's name (PI 30). The criteria of mastering this language do not involve 'a picture of the object com[ing] before the child's mind' (PI 6), they lie in the child's being able to carry out the task of fetching stones. This language game is expanded to include the builder directing the assistant on where to place the stones by pointing and saying 'this – there'; Wittgenstein asks whether the use of 'this – there' can be taught through ostension. He responds with 'What is supposed to show what they signify, if not the kind of use they have? And we have already described that' (PI 10). He likens the functioning of linguistic expressions to the use of tools from a tool-box, thereby emphasizing how linguistic expressions are used in a wide variety of contexts, and how the mastering of language is intertwined with the learning and practising of social practices. This leads to Wittgenstein stating one of the recurring themes of the *Philosophical Investigations* that 'the meaning of a word is its use in the language' (PI 43): people are confronted with the use of words, and the sense that they are able to make of these words is the meaning that those words acquire for them.

Wittgenstein goes on to show many different ways in which ostension can be used, e.g. pointing ‘to humans, to shapes, to colours, to pains, to moods, to numbers etc.’ (PI 26): once learnt, these ostensive definitions will go on to be used in very different ways and in situations that are not at all given in the actual act of being ostensively defined. He gives an example of someone learning how to play chess: in order for them to make sense of the definition ‘This is a king’, they would have to already know some of the rules of chess, the king’s ‘place’ being already prepared (PI 31). So, in contrast to Searle, who differentiates between games such as chess and speech acts, ‘The pieces... are not characteristically said to have a meaning, and furthermore when one makes a move one is not characteristically said to mean anything by that move’ (Searle, 1975), for Wittgenstein, making an utterance is similar to making a move in a game, since like in chess, which involves the projection of future moves and the comparison and evaluation of previous moves, understanding of a word is gained by comparison with ‘other contexts and by projecting its impact on future words and utterances’ (Duranti, 1997).

From this language game perspective, the criteria for determining correct or incorrect use of an expression reside not in reified conventions, but in how interlocutors actually go on to make use of that word in the broader language games in which the learning process is embedded.

Meaning as an internal state

Wittgenstein argues against the temptation of seeing meaning as residing in some internal state, showing how it could be invoked to explain the difference between pointing at an object and pointing at its colour, shape or number. As there is not one bodily difference between these different kinds of pointing, the temptation arises to explain the difference by saying that different mental acts accompany attending to a colour, shape, number, or other qualities. However, what at first seems a clear-cut simple act, namely concentrating on a colour, can be conducted in many different ways: comparing shades of blue when mixing paint or when looking at a blue signal-light (PI 33). What distinguishes these examples are the circumstances in which they occur; although they are superficially similar, each is situated within different language games. This difference does not reside in internal states, rather it is already ‘open to view’ in its surrounding activities.

Wittgenstein also argues against explaining that the meaning of (single-worded) commands, such as ‘Slab!’ in the language game in (PI 2) is a shortening or ellipsis of a sentence of the form ‘Bring me a slab’. This too is a temptation to explain meaning with internal processes: it sug-

gests that the speaker of ‘Slab!’ thinks the longer sentence in their mind whilst saying the shorter sentence. However, the speakers of the language in (PI 2) do not have the longer sentence as part of their repertoire. The similarities between the one word utterance and the longer sentence are due to their shared use within a language game of fetching bricks. Wittgenstein then homes in on the context surrounding a particular utterance, asking what makes a superficially ambiguous sentence such as ‘Five slabs!’ a request or a response. For Wittgenstein, the difference in meaning is not, as Searle posits, due to the speaker having internal states that correspond to formal propositional differences between requesting and responding. What sets requesting and responding apart is that they are embedded within observably different patterns of activity. The meaning of an utterance is socially situated in the language games that people play: the criteria for correct and incorrect use thus depend on public inspectable aspects of the relevant language game.

Rule-following and judgement

Despite the criticism of the use of inner or external objects as criteria that something has been used correctly (or not), the temptation still exists to explain the possibility of determining successful use as resting on rules that govern each instance of use, and that are ‘grasped in a flash’ (PI 138-139): ‘If meaning is a matter of rules of use, surely we ought to be able to state the rules’ (Searle, 1969). For Wittgenstein, whose approach is that of ‘re-arrang[ing]’ that which is ‘already open to view’ (PI 102) this search for objectified meaning is misguided, as it leads to inextricable paradoxes and contradictions. He gives an example of a projection of a cube deliberately chosen to result in a pyramid (PI 138), and asks how an internal mental state could determine the correct use of the concept ‘cube’, as the projection of the cube is susceptible to more than one interpretation (cf. Searle’s illocutionary intentions that include a representation of the conditions of satisfaction).

Wittgenstein turns to examining the state of ‘understanding’, giving descriptions of a child applying a mathematical rule of a series, showing how defining understanding as an internal event leads to the pursuit of a ‘chimera’, obscuring the entire learning process that is a gradual acculturation into a complex ‘form of life’. When this person says ‘I understand’ or ‘Now I know how to go on’ (PI 179), it only makes sense when spoken within a background of shared practices of learning mathematics: of importance is not some hidden representation in the pupil’s mind, but the publicly inspectable and verifiable continuation of the series as deemed correct by

the teacher, leading to Wittgenstein's exhortation 'Try not to think of understanding as a mental process at all' (PI 154). In the abnormal case, where the pupil says 'I understand' and continues with the series in a way that the teacher deems incorrect, Wittgenstein diagnoses a temptation, succumbed to by Searle, to explain meaning as a rule that anticipates all future applications of its use, showing that 'any interpretation still hangs in the air along with what it interprets' (PI 198a), resulting in a collapse in the distinction between correct and incorrect use that is brought about by the very rules that were supposed to guarantee this distinction.

Wittgenstein's solution is to look at rule-following from a different vantage point that does not attempt to endow rules with any privileged objectified status, a rule being a social practice that is followed before being formulated: a sign-post can function only by virtue of there being a social practice of using sign-posts, otherwise the sign-post would merely be a piece of wood or metal sticking out of the ground. The distinction between a correct and incorrect application is not due to anything intrinsic to the rule itself, rather it lies in the activities surrounding rule-following. Rules and rule-following cannot be artificially separated and abstracted from a community. The way out of the infinite regress is to recognize that 'there is a way of grasping a rule which is not an *interpretation* (PI 201), as when attempting to give reasons for carrying out a particular practice, 'my reasons will soon give out. And then I shall act without reasons. (PI 211). The regress halts at 'blind obedience' or conformity to 'bedrock' communal practices.

This highlights Wittgenstein's radical reformulation of language as rule-governed behaviour. Throughout the *Philosophical Investigations* Wittgenstein convincingly demonstrates that no rule or representation, whether external or even internal cognitive representation can function as guarantor of successful use. Further, the privilege traditionally ascribed to these rules leads to a collapse of the distinction between correct and incorrect use. For Wittgenstein, social practices are fundamentally prior to any speculation about which rule is or isn't being followed. Following a practice unreflectively simply 'is', since 'to think one is following a rule is not to follow the rule...' (PI 202). When confronted with harmonious practice, the temptation, succumbed to by Searle, arises to explain behaviour as rule-governed; the observed regularities of behaviour are presumed to be due to interlocutors drawing on rules that are 'grasped in a flash' (PI 191). Wittgenstein's solution is that these bedrock practices are customs that are followed unreflectively, giving members of a community a shared sense of the taken-for-granted: 'It is a custom that determines whether I am pointing in the direction of my fingertip or my wrist' (Teichman,

2006, pp. 567-579), and constitutive of these bedrock practices is how they are learnt.

Wittgenstein equates rule-following with the obeying of orders, and raises the question of who is right if two people react differently to the same orders or training: he imagines going as explorers to an unknown country, asking 'In what circumstances would you say that the people there gave orders, understood them, obeyed them, rebelled against them?' (PI 206). His answer is that this would only be possible by discovering the shared ways of speaking and acting, using the 'common behaviour' of mankind as a system of reference. This emphasizes that common to all language is a sense of harmony or agreement. He characterizes this with 'It is what human beings *say* that is true and false; and they agree in the *language* they use. That is not agreement in opinions but in form of life' (PI 241). In addition, for a communication system to be at all possible, there must also be agreement 'not only in definitions, but also in judgements'. This agreement in judgements that are bedrock communal practices is the foundation that provides the context in which objects can be used as standards or rules and also allows the distinction between correct and incorrect application.

Shifting the focus away from reified rules and meaning towards a view that rules can only function from within a community, means that to explain a rule is to show how it is learnt, i.e. how the novice is trained and acquires practical skills from someone who is adept at using them, thereby becoming acculturated into a particular form of life. This bypasses the problem of the need for the justification of bedrock practices. A novice, learning the alphabet is not required to justify that 'B' follows 'A', it 'just is'. The learning of bedrock practices of a community is the learning of the same sense of what is 'obvious'.

The dialogical turn provided by Wittgenstein changes the nature of the search for what constitutes successful communication, away from idealist explanations towards the realization that there is a whole host of very different language games of 'communicating successfully' that are embedded in even more diverse language games. Wittgenstein's recommendation is to focus on descriptive explanations of how people actually 'do' communicate successfully in real-world settings.

Perhaps the greatest difficulty with this account is that he does not directly address interlocutors' interpretations of that which is taken for granted in order to restore harmonious practice. However, even in this situation, it is clear that sufficient sharing of bedrock practices is a precondition of the possibility engaging with each other about correct or incorrect use. As noted,

this means attending closely to those situations in which people treat communication as unsuccessful. This is dealt with in more detail in the following section.

2.2.3 **Language game failure**

It is arguable that the notion of blind rule-following becomes problematic when explaining the actions of novices once they have become acculturated into their shared form of life and speak with a voice that is authoritative and representative of their speech community. Wittgenstein speaks of what 'I' do as a fully-fledged master of language (dropping the 'we'): the master's autonomy is grounded in a dependence on the community for his actions to be what they are, as the process of acculturation is constitutive of the meaning of his actions. Although implicit in Wittgenstein's treatment is the recognition that this acculturation is an ongoing process, incorporating (possibly novel) use in context, a tension emerges, as it is unclear how this acculturation changes when the novice becomes a master, since the master's autonomy is implicitly reliant on being as able to recognize mistaken use as any other master.

Throughout the *Philosophical Investigations*, the recurrent themes related to problems of understanding are associated with the ability of a novice to 'go on' and with philosophical problems relating to conceiving of understanding as an inner mental process. Wittgenstein's account is arguably less suited to instances where two fully-fledged members of a speech community disagree or experience difficulty in understanding, as his account of language acquisition rests on the novice-master asymmetry and the issues concerning mental phenomena are directed primarily at the theorist: 'our mistake is to look for an explanation where we ought to look at what happens as a "proto-phenomenon". That is where we ought to have said: this language game is played' (PI 656). This exhortation is of no assistance to fully-fledged members of a community where the nature of the game being played is precisely at issue, deeply problematizing 'how to go on'.

It is strange that Wittgenstein does not develop an account of how misunderstandings occur and are resolved (except as conceptual problems for the theorist), as from a Wittgenstinian viewpoint, it is precisely from within such games of dispute over meaning and validity that terms such as 'judgement' or 'validity' acquire their meaning, these games not only being restricted to the blindly obedient instructional master-novice dialogues, yet these situations remain underexamined.

Wittgenstein's emphasis on the harmonious practices of a community doesn't do justice to his

insights on the vagueness of the boundaries of language games (PI 68) or that '[there is] also the case where we play and make up the rules as we go along. And there is even one where we alter them—as we go along' (PI 83). He also recognizes that the boundaries between language games are both fluid and open to exploitation: 'To say "this combination of words makes no sense" excludes it from the sphere of language and thereby bounds the domain of language. But when one draws a boundary it may be for various kinds of reason. If I surround an area with a fence or a line, or otherwise, the purpose may be to prevent someone from getting in or getting out; but it may also be part of a game and the players be supposed, say, to jump over the boundary... So if I draw a boundary line that is not yet to say what it is for' (PI 499), yet these boundaries can cause ambiguities concerning whether a part is essential to a game (PI 567). In (PI 539) he gives an example of a face that can be taken variously as kind or malign, and suggests that its interpretation depends on probability and frequency of exposure to similar situations, as well as suggesting in (PI 652) that 'the hostile glance and the words may later prove to have been pretence'. This highlights the openness, contingency and revisability of meaning.

This appears to lead Wittgenstein towards the position that blind rule-following necessarily involves a plurality of interpretations that is brought about through the vagueness and openness in the rules as well as causing a change in the patterns of usage over time. As each member does not command a clear view over the full arena of language use, each member will have different frequencies of exposure (PI 539) to the local variants that constantly emerge and are modified through their use. This appears to be acknowledged by Wittgenstein who states that 'Misunderstandings concerning the use of words [can be] caused, among other things, by certain analogies between the forms of expression in different regions of language' (PI 90), although most of his solutions to problems of meaning are directed at the theorist who is still holding on to a misguided view of language. From within the language game, Wittgenstein's depiction is one of relative harmony, yet his account also suggests that the phenomenon of modification of meaning in use should present a recurrent co-ordination problem for interlocutors, requiring an account of how interlocutors resolve this problem of variation within a community 'as a matter of course'.

Where Wittgenstein covers problems of meaning between interlocutors, the resolution of the problem seems to rest on asymmetries between them: when saying 'When I gave you this rule I meant you to...' (PI 692) the criterion given for determining what was meant is whether the

speaker is a master of that 'technique'. In instances where neither interlocutor assumes the role of 'Master', he depicts the process as one in which they give each other justifications for their respective actions until they have reached bedrock 'and the spade is turned' (PI 217) whereupon they 'act without reasons' (PI 211) and say 'This is simply what I do'. After this, there is always the threat of violence which will lead them to 'act quickly, with perfect certainty, and the lack of reasons does not trouble [them]' (PI 212). This appears to conflict with his emphasis on the many different ways of 'going on' described above that interlocutors have at their disposal through exploiting the rules and vagueness of the language games.

(PI 206) provides an example of two people conversing who have been acculturated in different language games, involving an explorer going to a foreign country with a language that is completely unknown to the explorer. Wittgenstein asks 'in what circumstances would you say that the people there gave orders, understood them, obeyed them, rebelled against them...?', leaving the question open, commenting "The common behaviour of mankind is the system of reference by means of which we interpret an unknown language" and giving an account of how the explorer might teach the foreigner by 'expressions of agreement, rejection, expectation and encouragement'.

However, this problem of acculturation into different language games does not only apply to inter-community communication as in the example above, Wittgenstein's account suggests that interlocutors encounter similar difficulties within a language community and even within a particular language game, since if the rules can be made up by interlocutors as they go along, this necessarily entails the other interlocutors adopting this novel use if communication is to be successful. The question then arises as to how this is achieved? The examples provided by Wittgenstein do not translate well to the problem of two or more mature speakers within the same language community who experience problematic communication as it is unclear how two interlocutors who have both reached 'bedrock' and hence act differently, 'with right, but without justification' actually resolve the problem. The issue here is not that of finding or seeking to describe all the language games of miscommunication, as such an attempt is bound to founder on the same grounds as any attempt to describe all the different kinds of understanding, but of recognizing that the inherent openness of language, it not consisting of rails 'being laid out to infinity' (PI 218-219) simultaneously brings with it the issues of a plurality of interpretations and a multiplicity of resources that interlocutors have at their disposal in communicating, and

that interlocutors must necessarily draw on in order to resolve this inherent variation that exists within a community of language speakers.

2.2.4 Ethnomethodology

Ethnomethodology emerged out of a critique of the Parsonian account of social order, which proposed that members' subjective and individualistic aims and the ways in which they attempt to achieve them are reconciled with each other through socialization. Social cohesion is therefore provided by members sharing internalized norms that constrain members' aims and behaviour. The Parsonian account of 'voluntaristic action' (Heritage, 1984b) attempted to construct a scientific account that explained the motivation behind members' actions. For Parsons, these internalized norms and dispositions to act were not readily accessible to individual members, they were supposed to lie 'behind' members' behaviour, since exploitation of these norms was conceived of as threatening the social order by leading to 'unstable coalitions of unrest'.

Garfinkel incorporates many ideas from Schutz's phenomenological account in criticising Parson's treatment of members as 'cultural dopes' (Garfinkel, 1967). Schutz (1964) asserted the fundamental importance of including members' own constructs into any account of social order. Members stabilize objects as 'self-same' on a 'taken-for-granted' level, generally suspending doubt as to whether things are really as they seem and tacitly assume a 'reciprocity of perspectives' (Schutz, 1964). This allows the glossing over of different interaction histories, and different experiences. Garfinkel draws on Schutz's conception that 'common-sense' knowledge, rather than being static internalized norms that underwrote mutual-intelligibility, is a vague and unreliable 'patchwork' of beliefs, motivations and suppositions.

Garfinkel's 'ethnomethodologically indifferent' approach departs radically from Parson's. He focuses on how the members themselves actively draw on their assumptions in establishing a social reality that was 'identical for all practical purposes' (Heritage, 1984b). Far from being inaccessible to members, the cultural norms are used by members in making sense of and accounting for their everyday activities. Garfinkel emphasizes how the actions performed by interlocutors can contribute toward the definition of the social context. For example, formulating a turn as an answer can have the effect of making the preceding turn a question, potentially rendering the context into that of 'doing question-answering'. This example highlights the importance of the temporal unfolding of sequences of actions. Garfinkel's approach emphasizes how on a moment-by-moment basis members of a community reflexively shape and reshape the

communicative context in their communicative co-construction of social reality.

Thus, in contrast with the transmission model of communication, and in strong resonance with Wittgenstein, Garfinkel's solution to the problem of intersubjectivity is that it is in the first instance a practical problem that interlocutors face and resolve:

‘ “Shared agreement” refers to various social methods for accomplishing the member's recognition that something was said-according-to-a-rule and not the demonstrable matching of substantive matters.’ (Garfinkel, 1967, p. 30)

However, unlike Wittgenstein's reliance on imagined examples, the ethnomethodological approach is resolutely descriptive. Instead of constructing an artificial language game from a privileged theoretical position, the ethnomethodological approach is concerned with ascertaining the practices that members publicly orient to as such, from within the same culture as the analyst. It is these practices or methods and not shared beliefs or values which members draw on in order to make sense of their ‘form-of-life’. Importantly, these shared methods do not restrict members to a single shared view as the meaning of each message is actively constructed by interlocutors according to their individual experience of the present and previous interactions.

Interlocutors' actions are therefore made intelligible through methodical ‘common-sense’ reasoning that allows them to make sense of the collaboratively co-constructed social reality. A central plank in the ethnomethodological account of what underwrites this process is the notion of *accountability* of actions, which is used synonymously with ‘intelligibility’, in the sense that interlocutors are able to describe and therefore give an account of their actions and it is also invoked to describe the moral accountability that interlocutors face in being held accountable for their actions (Heritage, 1998).

Garfinkel demonstrates the impossibility of the ‘judgemental dope’ account of communication by showing how this would require endless detailing of what needs to be shared: in a series of experiments he assigned students the task of commenting on a transcript of a conversation they themselves had had. The commentary was to explain what they and the other parties to the conversation actually understood they had been talking about and were instructed to be concise, clear and distinct, and to explain ambiguities. In contrast to Wittgenstein's imagined scenario in which the ‘reasons give out’, reaching bedrock, the students found that each explanation introduced further ambiguities and complexities, complaining that they had to write ‘more and more’, with the task becoming increasingly Sisyphean as they attempted to clarify what had been said.

For Garfinkel, the possibility of communication rests on this communal and public practice of establishing what was said ‘was identical for all practical purposes’ (Heritage, 1984a, p. 58) which is accomplished through the glossing over of individual differences in following the assumption of reciprocity of perspectives. For Garfinkel, this highlights the importance of that which is not said, the tacitly held patchwork of vague knowledge and suppositions.

Against the view of the ‘judgemental dope’ requiring internalized rules that regulate behaviour and make understanding possible, Garfinkel states:

‘The big question is not whether actors understand each other or not. The fact is that they do understand each other, that they *will* understand each other, but the catch is that they will understand each other regardless of how they *would* be understood’ (Garfinkel, 1952, quoted in Heritage, 1984b, p. 119)

As evidenced by the experiment described above, what is ‘said’ in interaction stands against a backdrop of tacitly held assumptions about the nature of the unfolding interaction. These ‘seen-but-unnoticed’ methods of sense-making are brought into the open when interlocutors are confronted with a situation in which these expectations are breached.

To illustrate how these ‘seen-but-unnoticed’ practices underlie everyday social interaction, Heritage (1984b) gives the example of returning greetings: explaining why we return greetings is essentially a restatement of the norm (cf Wittgenstein), e.g. ‘I said hello to return the greetings’. Deviations from the norm, however, frequently involve justifications, excuses, such as claiming non-hearing, in which interlocutors account for their behaviour, yet even in these circumstances this orientation to the norm allows both interlocutors to make sense of the situation. Importantly, it is in these remedial actions, in the face of a threat to the social fabric, that interlocutors reveal their tacitly held orientations to the norm in a publicly accountable manner, rendering them available for the other participant, and hence the analyst.

The ethnomethodological program thus provides an empirical version of Wittgenstein’s insights into rule-governed behaviour, which unlike transmission models of communication, allows for the possibility of interaction without relying on sharedness. Understanding is seen as contingent, local, and always revisable, involving

‘...practices of etc., unless, let it pass, the pretense of agreeing, the use of sanctioned vagueness, the waiting for something later to happen which promises to clarify what

has gone before, the avoidance of monsters even when they occur and the borrowing of exceptions are all involved. I am proposing these as practices whereby persons make what they are doing happen as rule-analysable conduct' (Garfinkel, 1963).

These active methods employed by interlocutors facilitate the glossing over of individual differences in the playing of language games. Ethnomethodology offers an account that fills the gap outlined above, namely that the constantly evolving language games should lead to a greater threat of language game failure. If interlocutors are playing noticeably divergent language games, recourse to justifications of blind-obedience to a practice is of no assistance to interlocutors in situ. The ethnomethodological approach reframes this problem, providing analytical purchase on how interlocutors actually do in fact resolve difficulties in understanding, most notably through the use of conversation analytic repair, which is described in further detail in the next chapter.

Related to the threat of language game failure, Torode (2002) identifies a further gap filled by ethnomethodology: the recognition that games can be embedded within larger games, and that larger games can be built during interaction out of smaller ones leads to the problem of which language to select or play in any given situation. He argues that Wittgenstein seems not to acknowledge 'preview' games which allow interlocutors to collaboratively select which game is to be played next, thus pre-empting communicative difficulties caused by one interlocutor moving to a different game (such as Schegloff's 'Preliminaries to Preliminaries' (Schegloff, 1980) or the 'preface sequences' prior to story-telling (Sacks, 1992)).

2.3 Conclusions

The progressively more sophisticated (historical) confrontation of communication theory with communication practice highlights both the complexity of the process and the need for more systematic study. If language is the exchange of marks and noises we need to directly engage with understanding those processes of exchange, not their cognitive or realist correlates.

This chapters' juxtaposition of Wittgenstein's arguments with the conduit metaphor provides a powerful standard against which to measure further developments of analytic approaches to language, notably with respect to models that incorporate the notion of a 'shared' code. From the vantage point afforded by this chapter, these models, of which relevance theory is a prototypical example, pre-suppose from the outset that which is most in need of an explanation.

This chapter has shown how Wittgenstein's insights into rule-governed behaviour, including

his analysis of language learning as constitutive of co-ordination, demonstrate the need for a descriptive program that is in the first instance data-driven.

The next chapter will review psycholinguistic approaches to modelling communicative practices.

The progressively more sophisticated (historical) confrontation of communication theory with communication practice highlights both the complexity of the processes and the need for more systematic study. If language is the exchange of marks and noises we need to directly engage with understanding those processes of exchange, not their cognitive or realist correlates. The next chapter will review psycholinguistic approaches to modelling communicative practices.

Chapter 3

Empirical models of (mis)communication

3.1 Introduction

This chapter reviews empirical approaches to modelling dialogue that allow a characterisation of both successful and unsuccessful communication. The aim of this chapter is twofold. First, to review the range of empirical work that has been carried out on communication. Second, to discuss the relative emphasis in these studies on positive and negative characterisations of communicative success and the associated concepts of meaning. For completeness, empirical findings based on Grice's and Searle's intentionalist approach are included. However, the main focus of this chapter is on the collaborative model of Clark et al. (1996b) and the interactive alignment model of Pickering and Garrod (2004), and their accounts of the interactive processes involved in the development of co-ordination. Both models emphasize the importance of direct interaction but invoke different primary co-ordination mechanisms to account for how mutual-intelligibility is achieved and sustained. Of importance to both models is their incorporation of the conversation analytic notion of repair to describe how interlocutors resolve misunderstandings as part of this process.

3.2 Speech act processing

Although intentionalist models emphasize the role of the speaker in the generation of speech acts through prioritization of the speaker's intentions, the majority of empirical research in this field focuses on comprehension. Generally these experiments focus on the comprehension of

indirect speech acts in order to investigate the relationship between comprehension of literal and speaker meaning. The standard model used as a basis for these experiments is the ‘three-stage model’ (Krauss and Fussell, 1996) in which the hearer first processes the literal meaning and uses the literal meaning in conjunction with information from the communicative context as a basis for determining the speaker’s intentions with the assistance of conversational principles, such as Grice’s maxims. This model suggests that it should take longer to process indirect speech acts. So far, experimental results have proved inconclusive: studies on the comprehension of metaphor have shown that participants do not necessarily take longer to process than literal sentences (Gibbs, 1984) and are often arrived at even when they are irrelevant to actual use (Gibbs, 2002; Cacciari and Glucksberg, 1994).

The three-stage model has also been applied to the processing of indirect speech acts. Clark et al. (1975; 1979) conducted a series of experiments that investigated how comprehension of indirect speech acts is reflected in addressees’ responses. Clark identifies several features of indirect speech acts that contribute to their being understood as emphasizing the literal or the indirect meaning of the utterance. Broadly, these experiments addressed two kinds of influence on the interpretation of requests for information: properties of the individual speech acts and the effect of the communicative context. One such property is the ‘transparency’ of the request. Direct requests such as ‘I request you to tell me what time you close tonight’ explicitly mentions the ‘requestor’, the ‘requestee’ and the act that is being requested, thus the indirect version ‘Could you tell me what time you close tonight?’ is more transparent than ‘Stores seem to close early now’, which does not explicitly mention any of the three elements. Clark’s experiments involved making telephone enquiries about goods and services which varied in their level of directness. They were designed to elicit responses that could comprise two elements which corresponded to orientation toward the literal and the indirect meaning. For example, asking ‘Could you tell me what time you close tonight?’ can be taken to be both a question about the ability of the addressee to provide an answer and also as an indirect request for the closing time. Thus answering with ‘yes I can’, followed by ‘we close at 6’ responds to both, whereas responding with ‘we close at 6’ treats the direct question as *pro forma*. Clark observed that greater transparency in the request leads to responses that predominantly omit a response to the direct question.

Of importance to the concerns of this thesis are Clark’s findings that the ‘conventionality’ both of the form (idiomaticity) and ‘means’ (in the English language, indirect requests can be

performed *inter alia* by enquiring about the ability of the addressee to carry out the request, see Searle (1975) for a formalization in terms of preparatory conditions) had a similar effect on the response with greater ‘conventionality’ increasing the likelihood that the response would not include a turn-initial ‘yes’ or ‘yes I can’.

In section 2.1.2, one of the criticisms raised against intentionalist models is that the criteria guiding the inferential recovery of the intended meaning are perfectly robust and intersubjectively inscribed in the communicative context. Similarly, in the experiments described above, the addressee was strongly reliant on pre-existing conventions, although modulated by other factors such as transparency or implausibility of the indirect meaning. In addition, it is not clear how this speech act model could be used to address the development of co-ordination and potential changes in response patterns after repeated use of a particular request.

In a further set of experiments, Clark et al. (1983) investigated the role of the communicative context in interpretation, one of which involved showing participants a picture with Ronald Reagan and another less well-known politician, before asking either ‘You know who this man is, don’t you?’ or ‘Do you have any idea at all who this man is?’. The formulation of the first question suggests that the ‘man’ is known to the addressee, and in the second question unknown to both.

Although both questions are ambiguous, the patterns of responses indicate that addressees drew on their assumptions about mutual knowledge in interpreting the requests: 93% of respondents to the first question either selected Reagan or pointed at his photo, asking for clarification e.g. ‘This one?’; none of the respondents to the second question pointed at Reagan.

The experiments described above only involved dialogue consisting of two turns, highlighting the limitations of speech act theory (section 2.1.2) in accounting for the finer interactional contingencies existing within stretches of naturally occurring conversation.

Perhaps the strongest criticism of the application of intentionalist models to explaining actually occurring dialogue comes from anthropological studies by Rosaldo (1982) who conducted research on the Ilongots from the Philippines. In her study, Rosaldo criticizes Austin’s, Grice’s and Searle’s implicit assumptions about the importance of truthfulness and sincerity, as well as their emphasis on the role of intentions: the Ilongots do not have a social convention that is equivalent to the western act of promising that both Austin and Searle see requiring sincerity for successful performance; Rosaldo observes that the Ilongots were more pre-occupied with main-

taining social bonds than trying to ascertain the individual motives behind each other's intentions. Rosaldo's approach to the practices of the Ilongots thus echoes Wittgenstein's view of intentional action, describing it as embedded within a particular language game or form of life, rather than lying 'behind' the behaviour within an individual's mind:

An intention is embedded in its situation, in human customs and its institutions. If the technique of the game of chess did not exist, I could not intend to play a game of chess. In so far as I do intend the construction of a sentence in advance, that is made possible by the fact that I can speak the language in question. (PI 108).

Keenan (1974) also uses anthropological data in a study of Malagasy speakers to argue that Grice's CP is 'parochial and not universal' and that it can be seen as a quasi-logical specification of his own surrounding social norms. Keenan argues that Malagasy speakers do not adhere to the 'Maxim of Quantity': interlocutors regularly provide less information than is asked for, even when they have access to that information, as information that is known by other Malagasys is used as a means to gain prestige. In another view, it is possible to see that the Malagasys do indeed follow the 'Maxim of Quantity', as they can use it to derive their social status. However, as Kroch (1972) pointed out, Grice's CP can be shown to be sufficiently vague as to allow almost any proposition to be derived as an implicature from any other proposition. Although this does not invalidate the more general claim that interlocutors exploit each other's expectations, intentionalist speech act models do not address the main question of co-ordination in dialogue. A more direct empirical response to the issues raised in chapter 2 is provided by conversation analysis, which is described below.

3.3 Conversation analysis

Conversation analysis shares with speech act theory and Wittgenstein the assumption that language use is social action. As a development of the ethnomethodological program, its main aim is to describe the formal and structural mechanisms of conversation that interlocutors draw on in interaction. It shares ethnomethodology's concerns about the use of pre-conceived theoretical notions, and like ethnomethodology is concerned with mechanisms that interlocutors themselves can be seen to be orienting towards. So, in contrast to previous approaches that would classify social encounters according to the social groups of the interlocutors, such categories are only introduced in the analysis if the interlocutors make use of them during the interaction.

Conversation analysis takes as its main objects of inquiry transcripts of naturally occurring dialogue, eschewing the use of artificially invented examples and dialogue gained from experimental tasks that manipulate participants' behaviour, although increasingly, C.A. has been used to analyze the particularities of institutionalized talk (McHoul, 1978; Heritage, 1984a) which typically involve a restriction on the conversational mechanisms that interlocutors are able to draw on.

Unlike Chomskyan or Saussurean approaches, which characterize spoken dialogue as full of errors, leading to a focus on 'langue', the basic assumption of C.A. is that conversation is structurally organized, exhibiting 'organized patterns of stable, recurrent features' (Heritage, 1984b) which interlocutors are demonstrably oriented to. Unlike in speech act theory, the 'marginal', 'defective' cases are included in the analysis, as they can be used to delimit the scope of the particular conversational practices and serve as clear examples of interlocutors' orientation toward the norms that they breach.

C.A. pays close attention to the sequential unfolding of dialogue, with each turn normally displaying the understanding of the preceding turn, as well as projecting a relevant range of possible activities to be accomplished by the next speaker in the next turn. Sacks, Schegloff and Jefferson (1974) showed that the study of the succession of turns leads to the consideration of the succession of speakers which is accomplished with such precision by interlocutors that less than 5% of recorded speech is performed in overlap. Sacks et al. refer to this organization of orderly speaker exchange as the 'turn-taking system' which is not due to pre-allocation of turns, as in institutional settings. They propose that these seamless 'no gap - no overlap' transitions are due to interlocutors orienting to this system, identifying two components: the turn-constructural component and the turn-allocation component. The former is a minimally complete utterance, whether a single word, such as 'Hi' or a longer sentence such as 'So how are you doing?'. Importantly, the turn-constructural unit allows the addressee to 'project' a possible end of the unit, which constitutes a 'transition-relevance place', a juncture at which a speaker change could take place. The turn-allocation component is used to select the next speaker, who can be either the same speaker of the turn or a different one. Sacks et al. (p. 703) identify two rules that interlocutors orient to:

1. The speaker of a turn may either self-select and continue talking or select a different speaker who then continues.

2. The speaker of a turn doesn't select another speaker, in which case at the next transition relevant place either a new speaker self-selects or the speaker of the original turn continues talking.

This explains why overlaps occur at specific points, namely where the speaker of the next turn starts speaking in overlap at a projected end of the previous speaker's turn, as in the following excerpt where Curt's turn at (2) is taken as a simple question 'Who won?' : ¹

(Schegloff, 1987a, p.103)

- | | | | | |
|-------|--|-----|-------|------------------------------------|
| (3.1) | | (1) | Curt: | Well how were the races last night |
| | | (2) | Curt: | Who won [the feature] |
| | | (3) | Mike: | [Al won] |

Importantly, following the tenets of ethnomethodology, this also allows an identification and description of instances of deliberate interruptions, which do not (necessarily) occur at transition-relevance places, yet still evidence interlocutors' orientations to these rules in flouting them.

A consequence of focusing on the sequential unfolding of dialogue involving turns that project beyond the immediate turn leads to the consideration of units that are larger than the individual turn, the adjacency pair. This is a normative structure of two successive actions where the first action such as a question requires an immediately successive reciprocal second action in response (Goodwin and Heritage, 1990). Schegloff and Sacks (1973) provide the following formalization: (adapted from Levinson, 1983).

Adjacency pairs are sequences of two utterances that:

1. follow each other sequentially.
2. produced by different speakers.
3. ordered as a first part and a second part.
4. have a typology such that a particular first part requires a particular second (or range of second parts) e.g. offers require acceptances or rejections, greetings require greetings, etc.
5. having produced a first part of some pair, the current speaker must stop speaking and the next speaker must produce at that point a second part to the same pair.

¹The detailed transcription is reproduced in simplified form

The adjacency pair represents a significant advancement of the notion of speech as action. Although to a certain extent, Austin's requirement of 'uptake' by the addressee takes the reaction of the addressee into account, the adjacency pair does not prioritize the role of the speaker in the selection of the first pair part, nor attempt to give felicity conditions for its successful deployment. This kind of cross-turn interdependency allows the turn-by-turn monitoring of successive turns, with each turn displaying the understanding of the preceding turn and is a valuable resource for interlocutors in sustaining mutual-intelligibility. Importantly, the notion of adjacency is not hampered by the artificial distinction between illocutionary and perlocutionary effects of the first pair part. Instead, adjacency pairs are 'normatively' organized in that interlocutors are able to make use of this orientation to the norm in retroactively rendering the preceding turn as the first pair part, e.g. by formulating an utterance as an answer to a question. Orientation to this norm is particularly evident when an interlocutor does not produce a second pair part in response to the first, in which case the second pair part becomes 'noticeably absent' (Schegloff, 1968), as demonstrated by A's repeated attempts to 'pursue a response' (Pomerantz, 1984) in the excerpt below.

(Atkinson and Drew, 1979, p. 52)

- | | | |
|-------|--|--|
| (3.2) | | (1) A: Is there something bothering you or not?
(1.0 sec) |
| | | (2) A: Yes or no |
| | | (3) A: Eh? |
| | | (4) B: No |

However, the second pair parts are frequently preceded by embedded adjacency pairs that are completed prior to completion of the original, as in the following excerpt:

(Levinson, 1983)

- | | | |
|-------|--|--|
| (3.3) | | (1) A: May I have a bottle of Mich? (Question 1) |
| | | (2) A: Are you twenty one? (Question 2) |
| | | (3) A: No (Answer 2) |
| | | (4) B: No (Answer 1) |

Dialogue may contain many of these embedded insertion sequences (Schegloff, 1972). However, in such sequential trajectories where the two parts of the adjacency pair are separated by many turns, it is expected that the insertion sequences between the two parts of the original adjacency pair serve as 'preliminaries' that need to be resolved before returning to the second pair

part of the original. During the extended sequence, the second pair part is thus made ‘conditionally relevant’, a notion that replaces the stricter requirements of adjacency. The adjacency pair, which spans utterances and speakers, offers a subtly different treatment of co-ordination from speech act theory. While it would in principle to formalize relations between speech acts, Searle is primarily concerned with the illocutionary intentions guaranteeing felicitous performance of each individual speech act; perlocutionary effects are of secondary importance. By contrast ‘uptake’ and subsequent perlocutionary effects are intrinsic to C.A.’s methodology which draws on these effects in order to demonstrate interlocutors’ orientation to normative rules underlying co-ordination.

The notion of ‘preference organization’ builds up on this framework, based on the observation that while first pair turns such as apologies have a wide range of possible second pair types (acceptances, further accusations, challenges as to the sincerity of the apology, retraction of the accusation etc.) these second pair types are not treated equally by their speakers: ‘Dispreferred’ seconds are frequently accompanied by an account justifying the choice and why another alternative, the ‘preferred second’ was not performed, often being preceded with a hesitation and markers such as ‘uh...’, ‘well...’ or longer qualifiers such as ‘It’s not so clear at the moment, but...’ (Levinson, 1983).

3.3.1 **Repair**

‘There is a separate machinery designed for dealing with misunderstandings, and it draws attention to things that are not otherwise much focused on in conversation, apparently secondary bits of talk trying to “get things right”. Now we know that the “misunderstanding machinery” is itself “formal”, in the sense that it operates without regard to disagreements; it operates in lots of places where there is one or another sort of local failure of understanding’ (Sacks, 1992)

In contrast to models which seek to describe and formalize different kinds of error, this mechanism that was later developed into the notion of repair focuses primarily on the sequential unfolding of problems with the talk that are demonstrably oriented to by the participants themselves. Although the examples given in the literature on repair cover a wide range of ‘problems’, e.g. repairs on person reference, word-replacements, next-speaker selection (Schegloff et al., 1977), uncertain hearings (Schegloff, 2000), turn-beginnings due to overlap (Schegloff, 1987b),

rhetorical questions being mistaken for clarification requests (Schegloff, 1984), the kind of trouble is of secondary importance to the sequential nature of the interaction in which interlocutors identify, locate and correct the problem.

The mechanism of repair is an essential part of the conception that each turn displays the understanding of the preceding turn. It is the device used by interlocutors to signal and correct misunderstandings, although its use is not restricted to instances of miscommunication: as with all such devices, they can be exploited by interlocutors. Shifting the focus from objectively determined misunderstandings to misunderstandings or problems revealed by the interactants leads to repair covering a wider domain than that of correction, as it also covers instances where no discernible error is apparent to the analyst and also leads to excluding turns which an external observer might perceive as instances of miscommunication, but are not treated as such by the interactants. As mentioned above, the occurrence of a hearable error does not in any way entail that the other will initiate repair, as there is a preference for withholding reservations about prior talk until there is evidence to the contrary, following the ethnomethodological ‘etc.’ or ‘let-it-pass’ principle.

Schegloff et al. (1977) separate the processes of miscommunication into two main stages of repair-initiation and repair-outcome, noting that ‘the initiation of reparative segments and their completion (whether with success or with failure) can be quite distinct due to both stages being accomplishable either by the speaker of the trouble-source turn (henceforth TST) identified by the repair initiation or by the ‘other’ interlocutor, operating with a preference for both self-initiation and self-correction.

Schema adapted from (Levinson, 1983; Schegloff, 1992, 1997).

(3.4)	Turn 1:	There are two opportunities for self-initiated self-repair: 1) Immediately after the error; 2) At the end of the turn;
	Transition space:	Third opportunity for self-initiated self-repair
	Turn 2:	Fourth opportunity: either for other-repair or for other-initiation
	Turn 3:	Fifth opportunity: if other has initiated repair at turn 2, other-initiated self-repair
	Turn 4:	Sixth opportunity for repair: if B identifies through A’s response at (3) that B’s turn at (2) was predicated on a misunderstanding of (1), B has the opportunity for other-initiated repair

This schema describes the unfolding of the ‘repair-opportunity space’ (Schegloff, 1992) which results in an overwhelming preference for self-repair. Importantly, this is not to be thought of as cognitive preference. As described in section 3.3, the preference for self-repair is exhibited in deviations from this norm involving accounts (section 2.2.4) actions or involve additional interactional effort, and ‘are usually performed with delay within turns, and are variously softened and made indirect’ (Atkinson and Heritage, 1984, p. 53).

If the speaker of the TST forgoes the opportunity to repair their own utterance, whether they have signalled it as problematic or not, the other has the option of initiating repair and/or correcting the initial speaker’s utterance. Jefferson (1987) distinguishes between two types of other-correction: ‘exposed’ and ‘embedded’ correction. An example Jefferson gives of exposed correction is the following:

(Jefferson, 1987, p. 89)

- | | | | |
|-------|------|--------|--|
| | (1) | Meg: | It came from England Loren, |
| | (2) | Loren: | Ah-ah;;; |
| | (3) | Loren: | Ah-ah it’s stamped on the bottom |
| | (4) | Loren: | I:ndia,; |
| | (5) | Loren: | Madid in India |
| | (6) | Meg: | Ma:de. in India not madid, |
| | (7) | Loren: | .hh |
| (3.5) | (8) | Meg: | in India |
| | (9) | Loren: | (you said) m:madid. |
| | (10) | Meg: | You::shouldn’t say madid. |
| | (11) | Loren: | Madid in En:gland. |
| | (12) | Loren: | Da-uh i-in In:dia .hh |
| | (13) | Loren: | England. hh |
| | (14) | Loren: | They have s::ome, .hh they have them sort of like it in Indee(.)
in: England, .hh but it’s not exactly like it. |

Following Jefferson’s description of exposed correction, in (3) the correction of the term ‘madid’ becomes the interactional priority of the participants, giving rise to ‘accountings’ for the use of the term, in this case leading to the rejection of the proffered term ‘made’ and subsequent reuse of the original term ‘madid’. Jefferson distinguishes ‘exposed’ from embedded correction which occurs when the first speaker subsequently adopts the item / word used by the co-participant without discontinuing the ‘talk-in-progress’ or addressing the conduct of the other. Interestingly it is controversial whether embedded repair really constitutes repair precisely because it is not ‘displayed’ and is hence methodologically problematic for C.A.

(Jefferson, 1987, p. 92)

- (3.6) (1) Customer: Mm, the *wales* are wider apart than that.
 (2) Salesman: Okay, let me see if I can find one with wider threads
 (3) Salesman: (after looking through stock) How's this?
 (4) Customer: Nope, the *threads* are even wider than that.

The majority of other-initiated repairs do not perform correction of the trouble-source, even when 'other' knows the correction, producing instead 'a sequentially appropriate next turn on the basis of their understanding' (Schegloff et al., 1977, p. 378) of the previous turn, providing the speaker of the TST with the opportunity to self-correct. Unlike same-turn self-repairs which locate the trouble-source by repeating and modifying the problematic item and unlike embedded repairs that also implicitly locate the trouble-source by substitution, next turn repair initiation often involves the use of devices to locate the specific trouble-source in the previous turn.

Schegloff et al (1977) identify different kinds of other-initiation that vary in their strength at locating the 'repairable'. The weakest ('open-class' repairs, Drew 1997 such as 'Huh?' or 'What?' only specify that some difficulty was encountered with understanding the TST, but do not specify the nature of that trouble. Importantly, these kinds of repair-initiation are fully reliant on the sequential implicativeness of the talk: they are only usable in virtue of the fact that each turn displays the understanding of the preceding turn, since these open-class repairs do not specifically locate the trouble-source. Other alternatives, ordered in increasing strength of locatability are the following: questions, such as 'who', 'where', 'when'; partial repeats of the TST plus a question word, such as 'to the what?', 'next to where?', 'with whom?'; partial repeat of the trouble-source turn; a candidate understanding or confirmation check, e.g. 'You mean X?'. A preference for strong over weak locatability operates over the deployment of these initiation types that is apparent when the response to the initial other-initiation does not yield a satisfactory solution and leads to the subsequent deployment of another stronger repair initiator. This situation is one of the instances that can lead to other-initiated repair being deployed after the 'next turn position' following the trouble-source turn (Schegloff, 2000). Here too, preference is not to be understood cognitively. The normative preference for stronger initiation is hallmarked by 'weaker ones [being] self-interrupted in mid-production to be replaced by stronger ones' (Schegloff et al., 1977, p. 369).

Nearly all other-initiated repair is performed in the turn following the trouble-source turn. However, as mentioned above, other-initiated repair is occasionally deployed later due to unsuccessful repair outcomes that are then subsequently overridden by more specific repair initiators.

In these instances where other-initiated repair is 'displaced' from its naturally occurring position of 'next turn', the issue of how the repair initiation locates the trouble-source is of greater importance, as it cannot rely on being adjacent to the TST. This is made apparent by the absence of open-class repair initiators such as 'Huh?', 'Sorry?' or 'What?'. Instead, delayed other-initiated repair employs other techniques to locate the trouble-source, e.g. through repeating the trouble-source, repeating words that 'frame' the trouble-source or by using category-specific question words. Where open class repair initiators are employed, they are usually accompanied by a quotation of the trouble-source.

The notion of repair thus provides a radically different perspective on how successful communication is achieved. In contrast to speech act or intentional models, it is not presumed that misunderstandings are caused by some kind of mismatch of representations which prevents identical meaning recovery. Central to its ethnomethodological foundations is a focus on the procedural elements of how a particular turn comes to be oriented to as displaying a misunderstanding, and importantly, the mechanism of repair does not make any presuppositions about the 'meaning' of the utterances that is being communicated and subsequently does not require a specification or typology of kinds of misunderstanding that focuses on the source of the problem.

3.4 Clark and Marshall's model of definite reference repair

Clark and Marshall's (1981) model integrates the notion of repair into a framework which is primarily concerned with the failure to secure joint reference. Unlike C.A., which remains agnostic about the source of the problem, Clark and Marshall attempt to formalize the different kinds of knowledge and assumptions underlying successful communication in order to elucidate the necessary steps in order to resolve problematic reference. Clark and Marshall divide mu-

tual knowledge into four distinct types that have different associated assumptions (see table 3.1 below): Use of *community membership* as a basis for inferring mutual knowledge requires that interlocutors both believe that they mutually know they are members of a particular community (co-membership) and that they believe that all members of that community have that particular knowledge (universality of knowledge). It is supposed that ordinary communication involves deliberate use of specific terms to establish which communities the interlocutors hold in common

Basis for mutual knowledge	Auxiliary Assumptions
1. Community membership	Community comembership, universality of knowledge
2. Physical copresence	
Immediate	Simultaneity, attention, rationality
Potential	assumptions of 2. + locatability
Prior	assumptions of 2. + recallability
3. Linguistic copresence	
Potential	assumptions of 2 + locatability + understandability
Prior	assumptions of 2. + recallability + understandability
4. Indirect copresence	
Physical	assumptions of 2. + locatability /recallability + associativity
Linguistic	assumptions of 2. + locatability / recallability + understandability + associativity

Table 3.1: Clark and Marshall's 1981 model of definite reference repair

(Schegloff, 1972). Clark and Marshall conceive of mutual knowledge based on community membership as persisting over extended periods of time and being constantly renewed. They employ the metaphor of an encyclopedia with different chapters being associated with the knowledge of each community which is linked to a diary of situations involving triple-copresence. The second and strongest kind of evidence is based on *physical copresence* which is divided up temporally into the time periods during, before, and after the act of reference that have different associated assumptions. So a speaker's attempt to refer to an object that the hearer is not currently looking at, where both are in each other's field of vision, will rely on the potential for physical copresence and the assumption of locatability, giving rise to immediate physical copresence. Once they have looked at the object but are no longer attending to it is an instance of prior physical copresence requiring the stronger assumption that the hearer can remember the previous occasion of referring to the object (recallability). The third kind of evidence is *linguistic copresence* that occurs when an object referred to is absent from the communicative context. Here, linguistic copresence can only be prior or potential depending on whether the object referred to is established

before or after the moment of referring. Both prior and potential linguistic copresence require the same assumptions as prior and potential physical copresence, as well as the assumption of 'understandability'.

In contrast to community membership, mutual knowledge based on the other three kinds is conceived of as temporary, often being used on only one occasion. Also, the ordering from 2-4 is based on the increasing strength of the auxiliary assumptions.

As mentioned above, this model associates miscommunication primarily with the failure to secure joint reference. Clark and Marshall incorporate Schegloff et al's (1977) notion of repair, associating it with a strengthening of evidence in their copresence heuristics. They distinguish between two kinds of repair, horizontal and vertical. Horizontal repairs involve the provision of more information without altering the set of underlying assumptions.

(Clark and Marshall, 1981, p. 48)

- | | |
|-------|--|
| (3.7) | (1) A: I want that |
| | (2) B: which one? |
| | (3) A: The book right there on the second shelf |
| | (4) B: I still don't see which one |
| | (5) A: The green book on the second shelf from the bottom of that bookcase |

In the example above, both of A's repairs rely on potential physical copresence whilst increasing the specification of the reference, thereby strengthening the basis of mutual knowledge. When a referent is not physically available and a reference to an item in the preceding talk leads to problematic understanding, Clark and Marshall describe the process of disambiguation by the provision of further information that is also based on prior linguistic presence as leading to horizontal repair that in order to succeed 'must add or alter descriptors, not delete them'.

Problematic reference that relies on community membership is treated as a separate phenomenon that can be resolved by providing more information from the same chapter of the encyclopedia of community knowledge or by referring to a different entry which must either reinforce the belief that the interlocutors are members of the same community or increase the certainty that members of the community universally know the referent (see also Clark 1985).

Vertical repairs, by contrast, strengthen the kind of copresence in order to enable mutual knowledge to be inferred using weaker auxiliary assumptions. Clark and Marshall argue that repairs of utterances involving potential and prior physical copresence result in a shift to direct physical presence as it removes the need for the auxiliary assumptions of locatability and re-

callability, stating that it would be nonsensical for this shift to go in the opposite direction. An example of this shift is:

- (3.8) | (1) A: The book over there is mine
 | (2) B: Which one?
 | (3) A: (picking up book and showing it to B) This one

Arguably, this conception of vertical repair encounters problems within the level of linguistic copresence. While the examples above all describe the levels in terms of a determined shift from one to the other, they also describe the levels in terms of one level subsuming another. Thus, in the following example:

- (3.9) | (1) A: I think your idea is excellent
 | (2) B: What idea?
 | (3) A: A moment ago you mentioned going to a movie tonight

A's invocation of prior linguistic presence is described as an instance of immediate presence that 'brings back' the previous mention of an idea B was unable to recall. They reserve their discussion of these kinds of inconsistencies to community membership, stating that it cannot be ordered using the same criteria for strength as its associated auxiliary assumptions are different from those of linguistic and physical copresence and can be either stronger or weaker depending on the communicative context. These difficulties are used by Koschmann and LeBaron (2003) who use recordings from an operating theatre to argue that although the heuristics of copresence might appear to solve the conceptual problem of assessing mutual knowledge, when applied 'in situ' to genuine instances of interaction, the expectations of Clark and Marshall's model are not reflected in the dialogue, arguing that these conceptual difficulties reappear in assessing the maintenance of the auxiliary assumptions. (cf. Stahl 2002)

In summary, their conception of repair involves the strengthening of evidence for interlocutors' use of copresence heuristics by finding new evidence that requires weaker auxiliary assumptions.

3.5 Collaborative model

The definite reference model described in section 3.4 was mainly linguistic and was later transformed into the collaborative model (Clark and Wilkes-Gibbs, 1986; Clark and Schaefer, 1989; Clark and Brennan, 1991; Clark, 1996b) which subsumes physical and linguistic copresence as

well as the new notions of ‘joint actions’ and ‘joint projects’ into the personal common ground. Although Clark’s notion of joint project derives from Levinson (1992, p. 62): ‘...a fuzzy category whose focal members are goal-defined, socially constituted, bounded, events with constraints on participants, setting, and so on, but above all on the kinds of allowable contributions’, Clark’s rendition adapts Searle’s ‘we-intentions’ (Searle, 1990) and is rigorously defined as consisting of concise hierarchies, procedures and dynamics (cf. O’Connell and Kowal, 1995). The collaborative model (henceforth CM) retains the distinction between personal and communal common ground, which is treated as a reservoir of community-specific information (Clark, 1996a).

Being concerned with the joint construal of meaning, the securing of ‘uptake’ (cf. Austin) is essential to the grounding process, and is used to distinguish joint acts from individual acts. An often invoked metaphor is that of musicians playing a duet which requires both musicians to co-ordinate their pitch and tempo, as well as co-ordinating entry into and exit from sections that are further divided into phrases, measures and beats. Thus in the CM, an individual musician playing a single note practising on their own is not performing the same action as a musician playing the same note as part of an ensemble, since what is missing is the collaborative ‘we-intention’.

The CM conceives of dialogue as more than a succession of individualistic messages constructed by interlocutors: communication emerges out of the joint activity of participants interacting with the purpose of establishing that what has been said has also been understood (Clark and Brennan, 1991). This is inherently collaborative, requiring continuous co-ordination of both process and content, proceeding incrementally, building on and countering interlocutors’ initial assumptions, beliefs and knowledge in an iterative ‘grounding process’. Similarly to Allwood’s (1995) model, the CM (Clark, 1996b) proposes that collaborative acts can be divided into four distinct levels:

1. Executing and considering behaviour, e.g. through vocalization and attention.
2. Presentation and identification of a signal.
3. Signalling and recognition of a communicated meaning.
4. Proposal and consideration of joint projects.

Each level is to be thought of as connecting the actions of the participants, (for dyads) forming a single joint action. Importantly, understanding at each level requires evidence of understanding at a lower level.

The grounding process, which occurs at each level consists of iterative cycles of ‘presentation’ and ‘acceptance’ whereby interlocutors attempt to reach the grounding criterion (Clark and Schaefer, 1989), namely that they ‘mutually believe that the partners have understood what the contributor meant to a criterion sufficient for current purposes’. This process is triadic, in that an interlocutor’s acceptance of their partner’s utterance must be met by a further acceptance in order to be considered grounded. Successful completion of one such iterative cycle constitutes a ‘contribution’, after which both interlocutors mutually believe that the recipient of the original utterance has understood the original presentation. A key part to this conception is that utterances require ‘positive evidence’ of understanding, by asserting understanding, demonstrating understanding or displaying readiness to continue. Clark et al. (Clark and Schaefer, 1987, 1989; Clark, 1996b) order these demonstrations of understanding into a hierarchy of increasing evidence:

1. *Continued attention.* B shows he is continuing to attend and therefore is taken to be satisfied with A’s presentation.
2. *Initiation of the next relevant contribution.* B starts the next contribution that would be relevant at a level as high as the current one.
3. *Acknowledgement.* B nods or says ‘uh huh’, ‘yeah’ or the like.
4. *Demonstration.* B demonstrates all or part of what he has understood A to mean.²

These acceptance processes are recursive, constituting embedded presentations that require further ratification by the original speaker. What halts this recursion is the assumption that the evidence needed to accept each embedded contribution will decrease, until all that is required is the weakest evidence of continued attention.

The immediate criticism of this ‘Strength of Evidence Principle’ is that it excludes both contributions that are made unintentionally and communicative acts that do not elicit an appropriate response as in the following example:

adapted from (Allwood, 1992)

- | | | |
|--------|--|---|
| (3.10) | | (1) A: Think of a number |
| | | (2) B: (no response but hears, understands and thinks of a number) |
| | | (3) A: Double it and then multiply by three |
| | | (4) B: (no response but hears, understands and mentally carries out the operations) |

²A further level, ‘display’, consisting of a verbatim repeat of elements of the presentation was later abandoned due to it being one of the ways of initiating repair

Further criticism can be made of the evidence entailed by the ‘initiation of next relevant contribution’, as what is considered relevant is often precisely at issue for participants and must also be established collaboratively, as demonstrated by the participants in the studies of indirect reference described in section 3.2 above. More generally, it is unclear how to distinguish between a genuine acceptance and an unrelated new contribution; in the case of backchannel responses it can be difficult to tell whether the speaker’s continuing contribution is an acceptance of these backchannel responses or not. Similarly, not all interlocutors’ behaviours are necessarily part of the collaborative process of accumulating common ground: they make facial expressions, gestures, yawn, blush etc. and it is hard to tell when they should be considered part of the collaborative process.

3.5.1 Experimental evidence

The evidence for the CM comes from a series of experiments (Schober and Clark, 1989; Clark and Wilkes-Gibbs, 1986; Wilkes-Gibbs and Clark, 1992) that predominantly focus on referential, task-oriented dialogue. The prototypical setup involves assigning participants the task of arranging tangram figures in the same sequence; a screen that separates the participants requires them to collaboratively develop referring expressions for the figures. A key finding that consistently emerges in these experiments is that the length of the referring expressions used by participants decreases with successive mention. For example, the first attempts at reference are typically similar to the excerpt below:

(Wilkes-Gibbs and Clark, 1992, p. 184)

- (3.11) | (1) A: Okay the next one is... resembles someone that looks like they’re trying to
| climb stairs. There’s two feet, one is way above the other, and-
| (2) B: And there’s a, there’s a, a diamond on the right side, on a slant?
| (3) A: Yeah
| (4) B: Got it
| (5) B: Like, kind of looks like it’s off the back
| (6) B: Right, I got it

By contrast, after 6 trials, the referring expression is shorter and is completed in a minimal two-part contribution cycle:

(Wilkes-Gibbs and Clark, 1992)

- (3.12) | (1) A: Um, stair climber
 (2) B: Okay

To explain this contraction of referring expressions, Clark and Wilkes-Gibbs [p. 16](1986) propose the ‘principle of least collaborative effort’:

‘In conversation, the participants try to minimize their collaborative effort, the work that both do from the initiation of each contribution to its mutual acceptance.’ (Clark and Wilkes-Gibbs, 1986)

This contrasts with the ‘principle of least effort’ of individualist intentionalist models, as according to the latter, once participants have successfully established a referring expression, they should continue to use it without modification, since not to do so would violate the Gricean maxims of Quantity and Manner. Clark and Brennan (1991) argue that this Gricean misconception stems from the assumptions of unproblematic trouble-free presentations and acceptances that neglect phenomena such as speakers limiting the time they spend in planning and formulating utterances, leading to revisions, amendments or rejections. Similarly, utterances are frequently made without sufficient knowledge to make them error-free; instead interlocutors signal hesitancy about the sufficiency of their presentations with pauses, disfluencies or try-markers (as in the 2nd line of excerpt 3.12), potentially eliciting an alternative description from their audience.

Further evidence for this model is provided by experiments that focus on the levels of active participation in the dialogue. Experiments conducted by Schober and Clark (1989) and Wilkes-Gibbs and Clark (1992) investigated the co-ordination between dyads who were actively engaged in the tangram task and overhearers with different levels of participation. In Wilkes-Gibbs and Clark’s experiment, participants were assigned the roles of two different kinds of ‘overhearer’: a ‘passive bystander’ who silently observed the experiment, and an ‘omniscient bystander’ who observed the experiment via a video link in a different room. After six trials, both kinds of ‘overhearer’ attempted to match tangrams with members of the original dyad. Although both kinds of overhearer were exposed to the same utterances and consequently the same informational content, ‘omniscient bystanders’ took longer to complete the task and used longer referring expressions. Clark and Wilkes-Gibbs argue that ratification as a participant to a conversation is sufficient to increase communicative efficiency, due to mutual belief being established by participants’ awareness of each other’s exposure to the contribution cycles.

In Schober and Clark's (1989) experiments, participants were assigned to these two levels of participation with the added restriction of not being able to see the cards. They were instructed to attempt to replicate the same sequence of tangram figures being constructed by the dyad whilst overhearing their dialogue. It was found that both types of overhearer performed significantly worse than the original dyads, even when provided with a tape-recording of the conversation which they could pause and rewind to alleviate timing problems in switching between different tangrams. This suggests that interlocutors' opportunities for active participation, in particular their opportunities for actively clarifying possible ambiguities or misunderstandings is of key importance to the development of co-ordination.

This discussion of the effects of participatory status on co-ordination will be resumed in chapter 6.

3.5.2 Miscommunication in the collaborative model

The division between communication and what counts as miscommunication is clearly demarcated in the CM where it is conceived as an alternative pattern in the contribution cycle where positive evidence at a level appropriate to the last presentation (whether explicitly or implicitly through continued attention) is not demonstrated. In these instances, the CM invokes the conversation analytic notion of repair as a signalling device that requests help from the speaker to deal 'with a possible mishearing or misunderstanding of the contributor's presentation' (Clark and Schaefer, 1989). Although Clark et al. acknowledge the fundamental role that repair plays in the CM, they argue that as it only serves the purpose of signalling evidence of misunderstanding, it does not capture the 'positive steps' of mutual acceptance required in order to reach the grounding criterion and 'avoid trouble in the first place' (Clark and Schaefer, 1989). As mentioned in the previous chapter, ordinary transparent communication is seen as the norm, with miscommunication arising as its complication due to a mismatching within or between levels. According to the CM, during interaction, interlocutors believe that they are in one of the following four states:

1. *State 0.* B didn't notice that A uttered any e.
2. *State 1.* B noticed that A uttered some e (but wasn't in state 2).
3. *State 2.* B correctly heard e (but wasn't in state 3).

4. *State 3*. B understood what A meant by e.³

Successful communication ideally involves both interlocutors being in state 3 for each utterance and its constituent parts, which ordinarily involves understanding at each level presupposing understanding at the lower levels (downward evidence, upward completion and causality, Clark 1996b). Communication at level 3 constitutes successful communication that is demonstrated by the displays of evidence of understanding listed above. One form of miscommunication that occurs is when an interlocutor is in state 2 for at least one of the constituent parts of the preceding utterance. Here, the CM's incorporation of repair is used to index each of these four states by classifying repair initiations as either asserting or presupposing (Clark and Schaefer, 1989) the state of misunderstanding of the recipient of the problematic utterance. Clark and Schaefer (1987) subsume the notion of repair specificity (see section 3.3.1) into their model by combining it with the rule of 'least collaborative effort', proposing that in choosing a repair initiator, interlocutors follow the rule of selecting the strongest initiator that is consistent with 'understanding to a criterion sufficient for current purposes'. Thus the utterance 'I didn't hear the last word' is an assertion of being in state 1 and the utterance 'You saw what?' is a presupposition of state 2. In Clark and Schaefer's (1987) study of telephone miscommunication, they further subdivide the indices of (mis)understanding to include the assertion of no hearing, presupposition of no hearing, presupposition of incomplete hearing and presupposition of fallible hearing which are also ordered according to strength.

This incorporation of repair arguably leads to complications in determining the nature of an utterance: self-initiated self-repair and completions by an interlocutor are considered part of the presentation phase, whereas other-repair and other-initiated self-repair are considered part of the acceptance phase. Both can have embedded contributions in the form of clarification subdialogues, resulting in it being difficult to determine whether an other-initiated self-repair is part of the acceptance or presentation phase (Traum, 1999). Traum argues that this is of little use to an interlocutor in situ, as it is often necessary to look at both preceding and following elements of the conversation before determining the nature of an utterance. From a different view, however, this might be perceived as one of the merits of the CM, as it allows flexibility in interlocutors' co-construction of meaning and describes the emergent structure as negotiable and open to revision.

³This was formalized before the addition of joint projects to the CM. Presumably this would entail a further state of 'taking up the joint project proposed by A'.

In a sense, the CM is less specific than Clark and Marshall's model in detailing the reparative steps involved in re-establishing successful communication as it no longer explicitly describes the process of relaxing the auxiliary assumptions of copresence and temporality. In instances where an utterance does not secure reference, Clark et al. describe the ensuing embedded acceptance cycle involving the speaker refashioning their utterance through repair, expansion or replacement. These must in turn be ratified in order to be added to the common ground and can involve expressions of uncertainty via hedges, installments or try-markers. Thus, a plausible reconstruction of the CM's incorporation of repair is that it serves to identify the level of understanding that hasn't been reached thus guiding the response it engenders. However, where misunderstandings occur within the highest level, no mechanism is specified that might constrain the kinds of modification to the problematic turn made in response to the initiation of repair.

3.6 Interactive alignment model

The collaborative model outlined above assumes that interlocutors update their model of the other interlocutor on a turn-by-turn basis, incrementally adding beliefs established during interaction to the common ground. Although the CM adds a 'mentalist flavour' (Stahl, 2002) to the conversation analytic notions on which it builds, it doesn't provide a mechanistic account of production or comprehension in dialogue. It is this gap that the interactive alignment model (Pickering and Garrod, 2004) attempts to fill.

The main assumption of the interactive alignment model is that successful communication is underwritten by parity of representations between interlocutors. Pickering and Garrod argue that this is achieved via alignment occurring at the phonetic, lexical, syntactical and semantic levels, and occurs via a 'priming' process which at each level favours the most frequently and recently used representation.

The interactive alignment model (henceforth IM) has to a large extent been developed to account for dialogue phenomena that have been observed in a series of 'maze game task' experiments (Garrod et al., 1984; Garrod and Anderson, 1987; Anderson and Garrod, 1987; Garrod and Doherty, 1994). The basic experimental paradigm (explained in more detail in section 4.3.3) of the maze game task involves participants who are seated in separate rooms in front of a computer and communicate with each other via an audio link. The screen of each computer displays a simple maze consisting of interconnected nodes, through which participants must move their

respective position markers in order to reach a particular ‘goal’ node. Not all nodes are connected by paths, and some of the paths are blocked by gates which impede participants’ progress through the maze and must be opened through the use of switches. Participants are only able to see their own position marker, and although both participants’ mazes have the same underlying structure (configuration of interconnected nodes), the locations of their starting positions, goals, switches and gates are different. Importantly, each participant’s gates can only be opened by the other participant’s location marker being in the same node as a switch that is displayed on the *other* participant’s maze. This creates a recurrent co-ordination problem that requires participants to refer to the locations of these switches and gates in order to solve the maze.

The descriptions used to describe maze locations can be broadly classified into four types: Figural, which typically pick out a salient feature of the maze; Path, which involve tracing a route along the connections between nodes; Line descriptions that treat the maze as an array of horizontal or vertical (and occasionally diagonal) vectors; Matrix which refer to maze locations using Cartesian co-ordinates. Garrod et al (1984) argue that participants’ choice of these descriptions is not arbitrary, as in the paradigmatic tangram tasks of Clark et al. (Wilkes-Gibbs and Clark, 1992) but that they reflect the use of distinct situational models that represent the maze. These location descriptions are elaborated in more detail in section 5.4.1.

Evidence for this non-arbitrariness of descriptions is provided by Garrod and Anderson (1987) who observed that once a word had been used in association with a particular model, it tended not to be reused in conjunction with another model. For example, the word ‘row’ is frequently used to describe horizontal vectors in the maze, which can either be conceived of as an ordered set of elements as in ‘Switch is on the fifth row from top’ or as an unordered set: ‘I’m in the top row’. Garrod and Anderson observed that participants who had consistently used one of these types of line descriptions, when referring to a location using another type would introduce a novel term such as ‘line’, leading to the co-occurrence of descriptions such as ‘1st row’ and ‘top line’ but not ‘1st row’ and ‘1st line’.

In the dialogue as a whole, members of a dyad matched each other’s description schemes well above chance level: after 2 trials, 95% of interlocutors were recorded using the same scheme as their partner. Importantly, this convergence cannot be attributed to explicit negotiation as it only occurred 15 times in 56 games, and they occurred predominantly in later trials, associated with Matrix descriptions. In addition, schemes which had been introduced explicitly were as prone to

being responded to with a different description type as other ‘normally’ introduced descriptions.

Garrod and Anderson (1987) propose that the global alignment of participants in the maze game adopting the same referring expressions is achieved through a local and tacit process of ‘input-output matching’ (p. 207):

‘input/output co-ordination...may be simply stated as one of formulating your output according to the same principles of interpretation as those needed to interpret the most recent relevant input’

This basic mechanism has the effect of coupling the processes of comprehension and production both within and between speakers, and in the interactive alignment model, the reach of this mechanism is extended to operate on the phonetic, phonological, lexical, syntactic, semantic representations, as well as on the situation model. Between speakers, the input-output matching is conceived of as a largely automatic or ‘post-conscious’ (Bargh, 1989) process: a speaker’s utterance will activate (prime) the different levels of representation, thereby increasing the likelihood that these representations will be selected for reuse. Thus, although participants might initially start using different referring expressions, the most frequently used words, syntactic structures and their associated models will become increasingly likely to be reused, inhibiting the other competing referring expressions in a ‘winner-takes-all’ process. Thus stable, successful communication is conceived as a state of equilibrium in which both participants are sufficiently aligned to make sense of each other’s utterances.

In addition to alignment being driven by between-speaker priming operating on the same level of representations, alignment at one level of representation is conceived as leading to alignment at another through a percolation process, such that the use of a particular situation model will activate particular lexical items and vice-versa as in the ‘line’ vs. ‘row’ example above. In contrast to the CM, which treats the individual elements of referring expressions as essentially equivalent, this ‘percolation’ between levels emphasizes the development of a single coherent representation that places constraints on the kinds of referring expressions used, and allows an account of how participants refine the use of their referring expressions throughout the dialogue, developing local sub-languages for the task.

An additional contrast with the CM are the IM’s claims that alignment of representations ordinarily occurs without what they see as the ‘computationally intensive’ task of interlocutors modelling each other’s beliefs which is ‘an optional strategy that interlocutors employ only when

resources allow' (Pickering and Garrod, 2004). They explain the observations from Isaacs and Clark's (1987) experiment in which speakers appeared to tailor their utterances to their audience by stating that this kind of audience-design is a 'one-off decision' that is based on the 'perceived expertise of their addressees'. They contrast this with the dynamic inferential grounding process which they argue is implausible due to memory and processing limitations. Instead, interlocutors' convergence due to the automatic priming mechanism leads to the resultant shared information constituting what Pickering and Garrod call the 'implicit common ground'. The discrepancies that exist between interlocutors' models prior to interaction are reduced as the interaction proceeds by interlocutors making increasingly similar additions to their models by 'foregrounding the same information', resulting in implicit common ground approaching the content of full common ground.

3.6.1 Miscommunication and implicit common ground

The interactive alignment model describes miscommunication as arising due to inappropriately aligned representations that lead to a faulty implicit common ground. Pickering and Garrod argue that these faults are repaired by a mechanism that is ordinarily automatic and does not involve interlocutors making inferences about each other's beliefs. They distinguish this kind of repair that only draws on the implicit common ground from repair using the full common ground that is invoked when automatic repair is unsuccessful.

The interactive alignment model does not fully elaborate how repair is incorporated, it merely proposes that when an interlocutor encounters a problematic utterance (i.e. one that is checked against their representations), the interlocutor will reformulate the utterance in order to re-establish the implicit common ground. Although no mechanism is provided for selecting one formulation over another, Garrod and Doherty (1994) suggest that when interlocutors fail to secure reference with a particular referring expression, they adopt the previously most frequently used one.

Pickering and Garrod argue that the overlap of interlocutors' representations is ordinarily sufficiently large to obviate the need for interlocutors modelling each other's beliefs. When there is extensive misalignment of representations that cannot be resolved with implicit common ground or in dialogue that requires deliberate misalignment (deception, manipulation) the interactive alignment model allows for the 'resource-intensive' use of common ground that works on a level above the underlying egocentric alignment mechanism.

3.6.2 Emergence of conventions in linguistic communities

Although the CM recognizes that interlocutors develop shared referring expressions through interaction that are partner-specific, it is not well-suited to account for the dynamics of linguistic (sub)communities which can be seen in the distinction between personal and communal common ground. The former is treated as emerging out of interaction and negotiated afresh with each new interlocutor (Brennan and Clark, 1996), whereas the latter is depicted as a static entity that leads to clear binary distinctions of ‘inside’ and ‘outside’ information (Clark, 1996b) based on knowledge of community membership. This stems from the emphasis on dyadic communication that observes how interlocutors’ common ground develops out of individual interactions with members of their own community: ultimately these subcommunities are comprised of individuals and are sustained through interaction out of which emerge community-specific norms.

Compelling evidence for this process of emergence is provided by a series of experiments conducted by Garrod and Doherty (1994). Participants were assigned to one of three different groups: isolated pairs who always interacted with the same partner in subsequent games, a sub-community group whose members changed partners in each game, only interacting with members from the same sub-community, and a non-community group whose members always interacted with a new partner who was not drawn from the same community.

Although all groups developed community-specific referring expressions, they differed in the kinds of expression used and also in how this was achieved. Initially, pairs in the subcommunity group were less co-ordinated than the isolated pairs, using a wider variety of referring expressions and also locally matched fewer of each other’s turns. By the later trials, this pattern was reversed: participants in the sub-community group had converged on a single Matrix scheme and consistently matched each other’s expressions, in contrast to the isolated-pairs group where no increase in matching of expressions was observed. The non-community group initially started less co-ordinated than both groups, and became less co-ordinated on each successive game.

To account for these different patterns of co-ordination, Garrod and Doherty argue that the participants of the different groups use referring expressions that vary in their reliance on salience (Schelling, 1960), precedence and convention (Lewis, 1969): on starting the game, participants co-ordinate locally, relying both on salient features of the maze and precedence. In the isolated pairs group, the development of global conventions that are based on knowledge about the community is inhibited by any unsuccessful use of these nascent conventions making further use

unwarranted. By contrast, in the sub-community group, the group acts as a larger ‘repository’ of the conventions being developed such that on encountering individual difficulties in using a particular convention, the participants still maintain their belief that in the community as a whole the convention is being adhered to. Isolated pairs’ reliance on salience is substantiated by the observation that when their position markers were close to the goal, they used more ‘goal-related path descriptions’, overriding the use of convention. In addition, Garrod and Doherty argue that the sub-community achieves greater co-ordination from interacting with a wider variety of partners, drawing parallels with Posner and Keele’s (1968) account of the stability of a particular concept depending on the variety of exemplars that are used in its acquisition.

To explain the differences in co-ordination between the community sub-group and the group whose members interacted with different partners in each game, Garrod and Doherty claim that in the maze game this kind of stability is provided both by input-output co-ordination, and by resorting to the previously most commonly used expression when encountering communicative difficulties. Thus, in the subcommunity group, the common interaction history which develops with each game leads to the development of conventions that participants are able to ‘fall back on’ when encountering difficulties.

It is unclear to what extent the findings of this particular experiment are incorporated into the interactive alignment model: perhaps the greatest difficulty is accounting for Garrod and Doherty’s observation that in the initial trials, participants in the sub-community group predominantly use Line descriptions before converging in the later trials on Matrix descriptions. The conservatism of the interactive alignment model does not provide an account of how once a description has been established within a community, it might be supplanted by another one, yet a consistent pattern observed in previous maze game experiments (Garrod et al., 1984; Garrod and Anderson, 1987; Garrod and Doherty, 1994; Healey, 1997) is that across trials there is a general shift from ‘concrete’ (Figural and Path) descriptions towards more ‘abstract’ (Line and Matrix) descriptions which is not readily explained by initial frequency of use. Garrod (1999, pp 398-399) notes:

‘...co-ordinating on a common ‘description language’ was not simply a matter of sticking with the first reasonable scheme that emerged. It seemed to involve a much more extended history of development whereby speakers would explore different schemes in a co-ordinated fashion over a period of time. This is important because it

suggests a more complex underlying mechanism than one which simply sticks with the first mutually acceptable scheme encountered.'

It is unclear, however, how this could be integrated within the interactive alignment model, as its primary co-ordination mechanism is inherently conservative, operating on pre-existing representations. An alternative account for this development of descriptions is given by Garrod and Anderson (1987) who state that under certain conditions one interlocutor takes the lead, resulting in the 'follower' aligning on the 'leader's' representations, but not vice versa.

In the interactive alignment model, the primary mechanism underwriting semantic co-ordination is priming which ideally leads to all levels of interlocutors' representations being as identical as possible. When accounting for the higher co-ordination in the sub-community group, however, there is the implicit recognition that this relies on variation in the communicative context which will necessarily introduce more differences between individual interlocutors' representations. The same claim is made for dialogue both prior to and after convergence: the greater stability of Matrix descriptions in the sub-community group is explained by participants having encountered a greater variety of descriptions in the pre-convergence phase, and having encountered many different exemplars are able to develop more stable representations. In the post-convergence phase, stability is also explained by participants encountering the established convention in a wider variety of communicative contexts provided by different interlocutors, suggesting that the alignment mechanism is reliant on innovation for its guarantee of stability. This runs counter to the interactive alignment model's absolute emphasis on the priority of alignment of representations at all levels: the metaphor of two tightly-coupled production and comprehension systems is invoked as the paradigm case of successful co-ordination as it allows rapid matching of inputs and outputs, yet is revealed by these experiments to result in less development of co-ordination than within a group. Paradoxically, this implies that variation, and hence asymmetries between interlocutors' representations are intrinsic to both establishing and sustaining mutual-intelligibility.

3.7 Conclusions

Both the CM and IM do not readily address this global shift toward more abstract description types. The key problem as both Clark (1996b) and Pickering and Garrod (2004) emphasize, is that the selection of description types cannot be modeled as autonomous choices: interaction plays an essential role in co-ordination. However, their primary co-ordination mechanisms are,

in effect ‘semantically neutral’, focusing on co-ordination processes that are not sensitive to the different kinds of co-ordination implied by changes in description type. As a result, they operate in the same way regardless of whether participants are using Figural, Path, Line or Matrix descriptions.

This consistent finding that participants progressively develop more abstract descriptions as well as evidence provided by Healey (1997) that on encountering difficulties in co-ordinating, participants shift to more concrete (Figural or Path) description types indicate that participants’ choice between different descriptions and situation models is not neutral, suggesting strongly that the use of different description schemes index different levels of co-ordination.

The abstract schemes (Line and Matrix) invoke a relatively systematic and compositional model which abstracts away from each instance of the maze to an array of possible locations. The concrete descriptions (Figural and Path) depend much more on the concrete details of the layout in each case. For example, the use of abstract Line and Matrix descriptions allows participants to make direct comparisons within mazes (e.g. determining whether a location is adjacent to, or twice as high as another) from the description alone. This is not possible with concrete descriptions. Further, whereas Figural and Path descriptions refer to ‘missing’ boxes as ‘gaps’ or ‘holes’, in abstract descriptions they are integrated into the counts of rows or boxes used to specify a location. Boxes and gaps are thus given the same abstract ontological status and are combined into higher-order elements of rows or columns. As a result, co-ordination of abstract descriptions is more difficult since, by hypothesis, it involves co-ordinating on a semantic model that is not directly manifest in any particular instance of the maze.

These different location descriptions are also not equally entrainable in the same manner as the various referring expressions used in the standard tangram tasks: once a referring expression is associated with a particular tangram figure, the context permits that referring expression to be consistently used throughout the task. In the maze game, by contrast, concrete location descriptions such as ‘the sticking out bit’ are generally only reusable within the same maze, due to their dependence on individual salient features of that particular maze. When reused in a different maze, concrete location descriptions as a rule do not pick out the ‘same’ location, since in one maze the ‘sticking out bit’ might be on the left of the maze and on the right in another. This leads to considerable differences within concrete descriptions of opportunities for reuse, as well as differences in their stability across contexts: terms such as ‘middle’ or ‘top’ can be used in

nearly all mazes, referring to broadly similar locations. Unlike concrete descriptions, abstract descriptions, once successfully co-ordinated on, allow entrainment in a similar fashion to entrainment on tangram figures. The same abstract description refers (contingently) to the same location regardless of the individual features of the maze. This difference in the entrainability of descriptions provides further evidence that the choice between concrete and abstract descriptions is not arbitrary in Lewis's (1969) sense, and involves different levels of co-ordination.

The question then emerges of how this shift toward abstract description types is achieved by participants. Both models' emphasis of the importance of direct interaction are reinforced by experiments conducted by Healey (1997) which clearly demonstrated the emergence of different varieties of abstract description types in separate linguistic subcommunities, which were not readily mutually-intelligible between the different groups. As Garrod and Doherty (1994) note, across maze game experiments, individual dyads arrive at idiosyncratic schemes whose differences are not attributable to task experience alone.

However, both models do not satisfactorily address which interactive mechanisms are directly involved in this process. In the CM, it is assumed that interlocutors explicitly index beliefs which are added iteratively to the common ground, but no mechanism is provided that satisfactorily addresses how grounded referring expressions might be systematically supplanted by others. The strength of the IM over the CM, is that it sees the process of co-ordination as tacit, yet is also 'conservative' in that both alignment and the repair mechanism are presumed to operate only on pre-existing representations.

In summary, the maze game offers a means of directly investigating the development of semantic co-ordination. It is unclear from existing accounts how the shift toward abstract descriptions emerges out of interaction, but the compelling evidence that their use entails different levels of co-ordination allows their use as indices that permit the investigation of the individual mechanisms that underlie their emergence.

However, previous maze game studies have only indirectly manipulated the opportunities for direct interaction through the manipulation of group-composition that engender different individual interaction histories. What is desirable is the fine-grained experimental manipulation of the unfolding maze game interaction in order to establish which particular mechanisms are implicated in this process. The next chapter describes the experimental methodology adopted in this thesis which addresses this issue.

Chapter 4

General Methods

4.1 Introduction

Empirical investigation of dialogue phenomena has traditionally been restricted to tasks which do not allow fine-grained control over the unfolding interaction. Experimental manipulations typically involve relatively global manipulations of the communicative context, such as task-type, asymmetries in knowledge, levels of copresence or modality, and frequently require the use of a confederate in order to achieve the desired level of experimental control (Metzing and Brennan, 2003; Branigan et al., 2000; Boaz et al., 1998). In evaluating these tasks, it is not always clear to what extent the use of a confederate ‘contaminates’ the interaction due to adherence to a script preventing participants’ use of spontaneous dialogue.

Perhaps the most detailed studies of spontaneous dialogue have focused on descriptive analyses of natural conversation (e.g. Schegloff 1992) although the use of corpora is limited in that they provide only retrospective, correlational data that make it difficult to resolve conflicting interpretations of the phenomena.

Thus there is a large gap between the coarse-grained experimental techniques used to investigate dialogue phenomena and the fine-grained primary co-ordination mechanisms of the interactive alignment and collaborative models of communication which both emphasize the rich interactional work achieved by each turn.

What has so far been missing in current approaches to examining the development of co-ordination is the fine-grained equivalent of the ethnomethodological ‘breaching’ (Garfinkel, 1967)

	Participant A's view		Participant B's view
(1)	B: Obviously the relatives were coming like they do to see me		B: Obviously the relatives were coming like they do to see me
(2)		PROBE	A: relatives?
(3)		BLOCK	B: yeah
(4)		ACK	B: Ok
(5)	A: yeah		A: yeah

Table 4.1: Dialogue trajectory of an artificially generated clarification sequence. **PROBE** = probe-question, **BLOCK** = blocked response, **ACK** = acknowledgement.

experiments, namely the selective interference of particular aspects of the interactive mechanisms underlying co-ordination. The approach adopted here operates on the assumption that in blocking, disabling, or otherwise rendering the use of these mechanisms more difficult, interlocutors' attempts at maintaining mutual-intelligibility will allow better understanding of how these mechanisms function in spontaneous dialogue.

A pilot study was carried out to assess the viability of this approach, which is explained in more detail below.

4.2 Initial Pilot Study

The experimental technique adopted in this thesis draws on pilot studies carried out by Purver, Healey, King, Mills (2003) which involved developing a novel chat tool that facilitates context-sensitive fine-grained manipulation of turns during the unfolding of the text-based dialogue.

The basic experimental setup involved pairs of participants (dyads) seated in separate rooms, communicating using a synchronous chat tool. However, instead of turns being sent directly between chat clients, each turn is routed via a server which can systematically modify the turns in a variety of ways determined by the goals of the experiment.

The original pilot study used this experimental setup to investigate the interpretations given to clarification requests, using both a story-telling task (Bavelas et al., 1992) and a collaborative decision-making task. The experimental setup involved the chat server generating artificial clarification requests which took individual words from a participant's turn and incorporated them

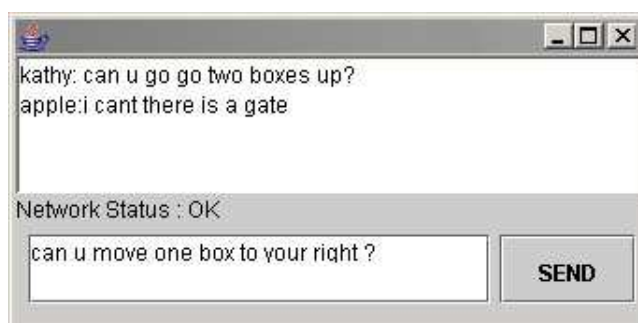


Figure 4.1: Client chat tool

into an artificial probe question which was sent back to the sender, yielding dialogue trajectories similar to table 4.1 above.

In Turns 2-4 the chat server engages B in an artificial clarification subdialogue, which blocks B's response before responding with an acknowledgement and resumption of the dialogue.

Importantly, debriefing after the experiment revealed that the interventions went undetected by participants, who responded to the artificially generated clarifications with similar proportions of answered/unanswered clarifications recorded in spoken corpora. Importantly, the results also showed that participants using the chat tool were sensitive to the development of co-ordination: participants' responses were affected by the level of grounding of the items being clarified.

This strongly suggests the viability of using text-based chat as participants' means of communicating in the maze game in order to perform fine-grained experimental manipulation of the unfolding dialogue.

4.3 Materials

4.3.1 Chat tool

This is a custom-built Java application similar to desktop messaging systems. The display is split into an upper window, a status bar and a lower window (see Fig. 4.1)

4.3.2 Chat server

The server is also programmed in Java, and in ordinary operation relays the chat turns between the chat clients. The server is also able to generate artificial turns that appear, to participants, to originate from each other or from a different source through modification of the name preceding the turn. In order to generate more convincing and 'natural' turns, it dynamically monitors the

participants' typing speeds and inter-turn delays.

When only one participant is engaged in an artificial dialogue, the server blocks text-entry in the other participant's chat tool, displaying an error message in the status window. To prevent association of the error messages with the interventions, the server sends error messages at random time intervals to both participants.

4.3.3 Maze game application

This custom-built Java application displays a simple maze consisting of a configuration of nodes that are connected by paths to form grid-like mazes (see Fig. 4.2).

The mazes are based on a 7x7 grid but with different configurations of nodes in each instance. Participants move position markers (a small black circle) from one node to another via the paths. Each move is recorded and relayed to the server where it is time-stamped and stored.

The game requires both participants to move their location markers from a starting location to a goal that is marked with a cross. Although the topology of the maze is the same for both participants, each participant has a different starting location, goal and marker, none of which are visible to the other. Some paths are blocked by gates (solid black lines) which can be opened by switches (grey coloured nodes). The locations of switches and gates are different for each participant. If a participant's location marker moves onto a switch on the *other's* screen, all the other participant's gates open. They close when the participant moves off the switch. This forces participants to use the chat tool to collaborate. In order for participant 1 to open his/her gates, he/she has to guide participant 2 onto a node that corresponds to a switch that is only visible on participant 1's screen.

Successful completion of each maze (when both participants' markers are on their respective goals) therefore requires participants to exchange descriptions of gate, switch, goal and position marker locations. Once both participants have reached their respective goals the next maze, with a new configuration, is automatically started. All three experiments involved participants playing 12 games.

The upper window displays the ongoing conversation, and the lower window is used for typing. Each turn is preceded by the name of the interlocutor. All keys pressed are recorded and relayed to the server, where they are time-stamped and stored. The status bar, a prominent single line of text that is controlled by the server, displays the activity of the other participant.

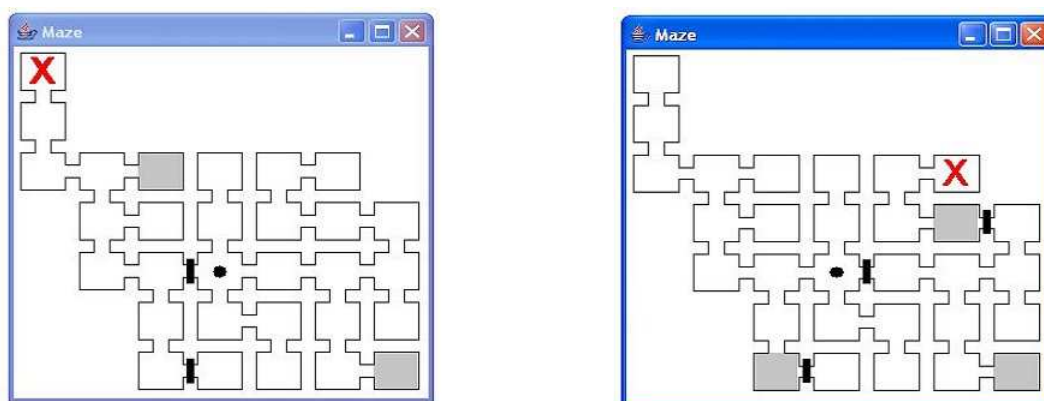


Figure 4.2: Example pair of maze configurations. The solid black circle shows the player's current position, the cross represents the goal point that the player must reach, solid bars the gates and grey squares the switch points

4.3.4 Procedure

Pairs of participants are seated in separate rooms in front of a desktop PC. On each PC a window containing the maze and a chat tool window are displayed. Participants are asked to select a nickname to be used in identifying chat turns and then wait for further instructions. Except when giving the initial verbal instructions the experimenter is seated at a third PC with screens to prevent any visual contact.

Participants are told that the experiment is investigating the effects of a novel chat tool on how people interact with each other. They are informed that their interaction will be recorded anonymously for subsequent analysis. Participants are also advised that they can request the log to be deleted and are free to leave at any time but would still receive payment in full. They are given a written description of the maze game and told that the experiment involves solving twelve mazes.

4.4 Comparison of media

One concern about the results from the chat tool experiments is whether they represent an artifact of differences between text and utterances as media or whether they bear on more basic aspects of the particular interaction mechanisms being investigated. An oft cited difference between spoken dialogue and text-based chat is the latter's 'quasi-synchrony' (Garcia and Jacobs, 1999): although the production and reception of messages are almost cotemporal (Clark and Brennan, 1991), in Clark and Brennan's terms, the messages are not fully simultaneous, since they are

formulated in a private window and are only transmitted when interlocutors press ‘enter’ / ‘send’. Unlike in spoken dialogue, there are no auditory or visual cues that allow the person formulating the message to simultaneously observe the recipient’s reaction, and the recipient is also unable to observe the formulation of the message. Any feedback can only be given on receipt of the full turn. Further, participants frequently only become aware that the other participant has been addressing them when they receive the full turn on their screen (Herring, 1999). This is not fully alleviated by the use of a ‘status bar’ indicator that shows whether the other interlocutor is typing, as this information provides no information that could be used to anticipate possible turn-completion.

In contrast to spoken dialogue, where the exchange of turns is locally managed through the anticipation of possible turn-completion points, in text-based chat, turn-allocation is also partly system-managed, giving rise to the phenomenon of *virtual simultaneity* (Garcia and Jacobs, 1999): turns that are sent at exactly the same time are ordered sequentially in the same manner as turns between which there is a large gap. Whilst the chat-text displayed on each user’s screen appears to follow the ‘no gap, no overlap’ rule (Herring, 1999) this rule is effectively violated by actual turn formulation being frequently performed in overlap. The privacy of turn-formulation can lead to interlocutors being unaware of whether the other is responding or not, leading to a further turn being sent before receiving the response to the first (Herring, 1999) Further, on receiving the response, it is not always apparent whether the turn is in actual fact a response to the same turn. Garcia and Jacobs (1998) argue that this phenomenon of *phantom adjacency* is akin to Garfinkel’s (1967) computer therapy experiment, in which turns that appear to be the second pair parts of some prior adjacency pair are interpreted as such, even though they are responding to some other turn. In this sense, the quasi-synchrony of text-based chat introduces more slippage into the dialogue. Thus, if not paying close attention to the unfolding dialogue, it is not evident from the chat window what the time-interval is between two successive turns, providing fewer cues as to whether the following turn is actually a response to the prior. Even if there is sufficient delay that is also apparent to interlocutors, there is no guarantee that the turn is in fact a response, since constraints of the exchange mechanism lead to messages being ordered according to when they were received, ignoring turn-initiation times.

These phenomena described above all have the effect of interfering with the sequential coherence in the dialogue (Herring, 1999). In spoken dialogue there is the expectation that turns

that belong together should be adjacent: a request should ordinarily be followed by acceptance or compliance; if missing, the turn that replaces it is understood to be a preliminary to completing the response. Herring argues that in text-based chat, ‘violations of sequential coherence are the rule, rather than the exception’. In a study of text-based chat, Herring (1997) observed the frequent occurrence of physically adjacent turns that were unrelated to each other, even turns that were intended to be related violated the Gricean maxim of relevance in 33% of the cases. The upshot of this is that in text-based dialogue, relations between turns appear harder to track than in spoken dialogue, not only making cross-turn anaphora more problematic, but also exacerbating the problem of establishing the relevance between successive turns.

These reservations notwithstanding, even if the dialogue phenomena observed were found to be primarily artifacts of the media, this would not detract from their significance, due to the paucity of experimental studies that directly address the development of semantic co-ordination in text-based CMC. In addition, the empirical investigations described in the ensuing chapters provide an important counterpoint to existing theories of spoken dialogue. It remains to be seen exactly where the differences lie. In order to address these concerns, the results from the experiments are assessed against the original maze game studies (which also involved constraints on participants’ communication via the use of an audio link).

Chapter 5

The effect of sequential coherence on semantic co-ordination

5.1 Introduction

As discussed in section 3.5, the collaborative model requires positive evidence of understanding to be provided turn-by-turn in order for information to be added to the common ground. Unlike in monologic accounts, turns cannot be understood when divorced from their communicative context: even utterances as straightforward as ‘Two right from the top left hand square’ are underdetermined until it is known whether they are the first or second part of a request, question etc., a process which is assisted through interlocutors providing each other with positive and negative evidence about the intelligibility of each other’s utterances. The interactive alignment model also strongly emphasizes the role of feedback in allowing the convergence of mental representations through the coupling of the production and comprehension processes.

The experiment described in this chapter seeks to interfere with these turn-by-turn displays of understanding to reveal the mechanisms involved in the development of semantic co-ordination by prohibiting the interleaving of participants’ chat-text.

5.2 Sequential coherence

The resources that interlocutors draw on to provide turn-by-turn evidence of understanding rely strongly on the local context in order to be understood: in spoken dialogue, the sequential coherence provided by the normative turn-taking rules and conditional relevance (see section 3.3)

are of prime importance to interlocutors in maintaining mutual-intelligibility. In particular, conversational mechanisms such as utterance-anaphora and simple feedback morphemes (Allwood, 1992), e.g. ‘Ok’ or ‘Huh?’ require for their successful deployment a high level of inter-speaker coherence which allows the immediate local context to be used in determining the other pair part. In the case of clarification requests, the first pair part is usually in the immediately preceding turn (Purver et al., 2003).

Although both the CM’s and IM’s primary co-ordination mechanisms emphasize the importance of feedback, neither model gives an account of which particular mechanisms are essential to the development of semantic co-ordination. The experiments conducted by Schober and Clark (1989) demonstrated the importance of active participation in co-ordination, but the main criteria used to assess differences in co-ordination were performance measures of timing and correctness of the tangram sequences.

While the studies of text based chat described in section 4.4 were primarily concerned with enhancing coherence in dialogue, the observations suggest that text-based chat can be manipulated to become less sequentially coherent. The research closest to the method adopted here was conducted by Anderson, Beard and Walther (2000) who used a chat tool that prohibited interleaving of chat-text by separating each participant’s text into different windows. They observed that participants stopped relying on turn-adjacency for understanding each other’s turns and adopted, instead, a ‘typing–overlap’ strategy of alternating between typing multiple turns without reading the unfolding dialogue and pausing to read the elapsed turns.

What might be the effect of reducing sequential coherence on the development of co-ordination in the maze game? Since text-based chat provides no paralinguistic cues, all positive and negative evidence of understanding must be gleaned from the text in the chat window. Although the collaborative model does not directly address sequential coherence in dialogue, a plausible reconstruction would predict that a decrease in sequential coherence should result in longer contribution cycles due to the greater effort required for continuing with the ‘next relevant contribution’, since what is considered relevant and what is considered to be the next contribution is highly dependent on sequential coherence. More importantly, the use of both positive and negative feedback should be more ‘costly’ (Clark and Brennan, 1991), due to the greater difficulty in drawing on the immediate local context, resulting in more explicit forms of feedback and/or explicit displays of understanding.

The interactive alignment model also strongly emphasizes the role of feedback in the convergence of interlocutors' mental representations. Reducing the sequential coherence should increase the difficulty of interlocutors' attempts to establish the relationship between turns, hindering the automatic process of input-output matching.

It is unclear, however, how to distinguish empirically between both models' incorporation of repair. The IM sees automatic repair using implicit common ground as the norm, yet provides no clear criteria for when miscommunication leads to drawing on the full common ground. A further difficulty is identified by Barr et al. (2004) who argues that Clark's formulation is equivalent to implicit common ground, rendering the distinction vacuous.

In both models, reducing sequential coherence should lead to more instances of misunderstanding, due to participants being prevented from relying on turn-adjacency and relevance. These misunderstandings should require greater effort in their resolution due to the reduced effectiveness of repair in 'locating the repairable' (Schegloff, 1992).

Both models, however, do not provide an account that is sensitive to the kinds of descriptions being used as they are semantically neutral. By hypothesis, decreasing the sequential coherence of the dialogue should lead to a decrease in participants' ability to manipulate each other's turns, and consequently lead to greater difficulties in establishing sufficient co-ordination to enable the use of abstract descriptions.

5.2.1 **Experimental hypotheses**

1. The collaborative model predicts that the decrease in sequential coherence should lead to greater difficulty in completing the contribution cycles, resulting in less contracted referring expressions in dual window dialogue.
2. All things being equal, the interactive alignment model predicts that interfering with the sequential coherence of the dialogue should result in weaker priming, and therefore less alignment, since participants will find it harder to establish the first part of the adjacency pair of which the participant's turn is the second part. Further, it predicts that the increased effort required to resolve misunderstandings should also weaken alignment.
3. The abstract schemes (Line/Matrix) should be used more in the single window condition than in the dual window condition.



Figure 5.1: Dual window chat tool. Each half only displays text from one participant

5.3 Methods

The maze game application and the chat server are the same as described in section 4.3. In this experiment no artificial probe turns were generated by the server.

5.3.1 Chat tool

Two different chat tool implementations were used. In the baseline condition, the same version of the chat tool as described in section 4.3.1 was used. In the second condition, participants used a variation of the chat tool¹ similar to the two-way chat system of Anderson et al. (2000) which separates the chat-text into two separate windows that only display chat-text from a particular participant (see Figure 5.1). Thus, in the dual window condition, after formulating a turn and pressing 'enter' or 'send', the new turn appears in the same window as all other turns by that particular participant and the other window only displays turns from the other. This dual window setup prohibits any interleaving between the turns of different participants.

5.3.2 Participants

23 pairs of participants were recruited, 29 male and 17 female, from undergraduate students. They were recruited in pairs to ensure that they were familiar with each other. This was for practical reasons, due to the high cancellation rate (>30%). Importantly, this increase in familiarity should, all things being equal, decrease the level of semantic co-ordination developing during the

¹The dual window modification of the chat tool and the collection of data were carried out by Shayan Akhgar as part of an undergraduate project.

experiment, thereby reducing the possibility of Type 1 error. Only participants who had some previous experience of using internet chat software such as ICQ or Microsoft Messenger were selected for the experiment. Each participant was paid £7.50 for participating in the experiment. Two pairs were replaced because they failed to meet the criterion of native competence in English, which was having one or more native English-speaking parent or having grown up in an English-speaking country.

5.3.3 Procedure

Participant pairs were randomly assigned to one of two conditions: the baseline condition, using the chat tool described in section 4.3.1, and the dual window condition, in which participants used the modified chat tool which separates text from both participants into separate windows. Twelve mazes were presented in random order to each pair. Overall, this resulted in Window type as between-subjects independent variable and Experience as within-subjects independent variable.

5.4 Results

Overall, 8036 turns were classified according to the criteria developed by Garrod et al (Garrod et al., 1984; Garrod and Anderson, 1987) which classifies location descriptions into Figural, Path, Line and Matrix, corresponding to different underlying mental models of the maze. These four description types are illustrated below:

5.4.1 Description types

Figural: a heterogeneous category of relatively concrete descriptions that draw on some specific element of the overall configuration or distribution of particular features to identify a target location.

- (5.1) | (1) A: You see the sticking out bit at the top?
 (2) B: Yep
 (3) A: I'm on the bottom right one and the switch is right above it

Path: involves identifying a route to be traversed through the maze to the target location. Typically they involve identifying a salient feature of the maze that is close to the target that serves as the starting point for the traversal and are sensitive to the specific layout of boxes and connections in the maze.

- (5.2) | (1) A: Where are you now?
 (2) B: See where your switch is?
 (3) A: Yep
 (4) B: Go up 1, 2 right, 1 down

Line: classifies the maze into a set of line elements corresponding to rows, columns or diagonals.

Line descriptions identify both the target line and the target box as a position along it.

- (5.3) | (1) A: Where's your switch?
 (2) B: It's in the bottom box in the second column from the right.

- (5.4) | (1) A: Am close to goal, where are you?
 (2) B: I'm in the third row, fifth to the left

Matrix: introduces a Cartesian co-ordinate system with locations identified via the specification of two vectors either as rows and columns or in terms of numbers or letters for each axis.

- (5.5) | (1) A: My switches are at 4,6 5,4 and I'm on 3,4

- (5.6) | (1) A: My goal is b2

- (5.7) | (1) A: I need to get to 2nd row 5th column

Coding the turns with Garrod et al's criteria gave rise to difficulties that are specific to text-based chat. Due to each turn only being sent when participants press 'Send' or 'Return', participants frequently split their contributions into multiple turns. Thus it is often unclear whether to classify the extended sequence as a single location description or as separate reformulations, in particular when interleaved with chat-text from the other participant. For consistency, it was decided to code ambiguous sequences as belonging to the same location description unless participants explicitly treated the descriptions as different.

- (5.8) | (1) A: I'm on the leftmost square in
 (2) A: the fifth line

- (5.9) | (1) A: Switch is on the sticking out bit
 (2) A: That's actually the 3rd column, second row

According to the schema adopted in this thesis, excerpt 5.8 is classified as a single Line description, and excerpt 5.9 as a Figural and a Matrix description.

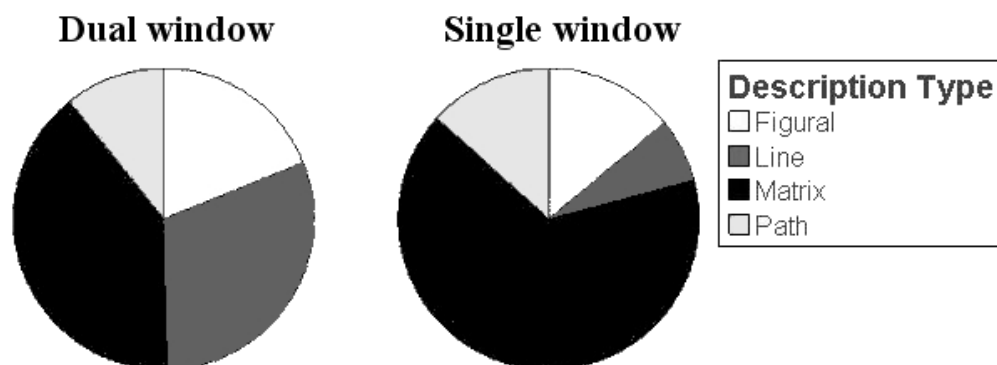


Figure 5.2: Global distribution of description types in dual and single window dialogue

Figure 5.2 shows the distribution of description types in the baseline and dual window conditions². The global pattern in the baseline condition is comparable to the ‘isolated pairs’ group in Garrod and Doherty’s (1994) experiment that reported a distribution of 10% Figural, 13% Path, 7% Line and 66% Matrix.

Comparing the overall distribution of description types in the baseline and dual window condition yielded a reliable difference (Multinomial Regression $Chi^2_{(3)} = 279.17$, $p = 0.00$), showing that the use of the dual window significantly affects co-ordination in that group. Although Matrix descriptions predominate in both conditions, they are more prevalent in the single window condition.

To provide focused tests of the development of co-ordination, four additional comparisons using Multinomial Regression were carried out. Firstly, the distribution of description types used in the first four games (early) was compared between conditions. The difference in use of description types is reliable (Multinomial Regression $Chi^2_{(1)} = 5.9$, $p = 0.015$), with participants in the dual window condition using more abstract (Line/Matrix) descriptions than in the single window condition. In the last four games (late), this pattern is reversed, with participants in the dual window condition using significantly more concrete (Figural/Path) descriptions than the baseline group $Chi^2_{(1)} = 5.73$, $p = 0.017$).

Focusing on within-group development of semantic co-ordination, the distribution of description types in the first four games (early) of dual window dialogue was compared with those used in the last four games (late). This showed no reliable difference $Chi^2_{(3)} = 6.51$, $p = 0.09$). The same comparison between description types in early and late games was made in the baseline

²Pie-charts were used because these are compositional data




Figure 5.3: Distribution of description types in dual and single window dialogue. Further subdivided into the first four games (early) and last four games (late).

condition. This showed a reliable effect of experience $Chi^2_{(3)} = 58.3, p = 0.00$). This pattern is illustrated in Figure 5.3.

Entrainment

Following Garrod and Doherty (1994) the degree of entrainment between participants was calculated by comparing each location description with the previous location description used by the other participant. If both were of the same type, an entrainment score of 1 was recorded, whereas if both were of different types, a score of 0 was recorded for that particular turn. For each game, each participant's score was added up and divided by the number of turns involving location descriptions made by that participant in that game. If participants consistently match each other's location descriptions the entrainment score would be a maximum of 1. Conversely, if participants consistently use different location descriptions the entrainment score would be 0. Importantly, this score is participant-specific, since it is conceivable that one participant might always match the other's description, whilst the other consistently uses a different type

Figure 5.4 shows the development over time of the entrainment scores of participants in both

conditions. It is immediately apparent that in the dual window condition, participants are more entrained than in the single window condition. This is confirmed by subjecting the scores to a one way analysis of variance, with Window type (dual/single) as between-subjects factor ($F_{(1,150)} = 7.11, p = 0.008$). This confounds the basic predictions of the interactive alignment model which would expect to see greater alignment in the baseline condition due to lower sequential coherence resulting in weaker priming.

Interestingly while the entrainment scores in both conditions are initially low and increase rapidly for five games, the entrainment scores taper off in the second half of the experiment. This result is comparable to the observations of Garrod and Doherty (1994) who also found a ‘levelling off’ in co-ordination in the ‘isolated pairs’ group toward the end of the experiment.

5.4.2 Adaptation to the media

The log files provided four measures of task performance. Firstly, the maze completion times for each of the 12 games was calculated from the log files and analysed in a single factor analysis of variance with Window type as fixed factor. There was no evidence of a reliable difference between groups. ($F_{(1,274)} = 0.7, p = 0.4$). In the dual window, participants took on average 3:40 mins per game, and in the single window participants took 3:20 mins.

The second measure of task performance was the number of turns produced by each participant. This was analysed in a one-way analysis of variance with Window type as fixed factor. There was no reliable effect of Window type ($F_{(1,274)} = 0.158, p = 0.691$). Dual window participants formulated on average 28.4 turns per game, and in the single window condition, participants formulated 29.8 turns per game.

The third measure was the typing speed of participants’ turns. An analysis of variance with Window type (dual/single) as fixed factor and game number as random factor showed a reliable main effect of Window type ($F_{(1,11)} = 7.22, p = 0.021$), no main effect of game number ($F_{(11,11)} = 1.82, p = 0.17$) and no interaction ($F_{(11,528)} = 0.301, p = 0.99$). In the dual window condition participants’ mean typing speed was 4.2 characters per second, and 3.8 characters per second in the single window condition.

The fourth measure was the change in turn length with task experience. A single factor analysis of participants’ turns in the baseline condition with Experience (early/middle/late) as within-subjects random factor showed a reliable main effect of Experience ($F_{(2,4283)} = 50.9, p = 0.00$) and a reliable reduction in turn-length (Linear trend, $p = 0.00$). The parallel analysis of

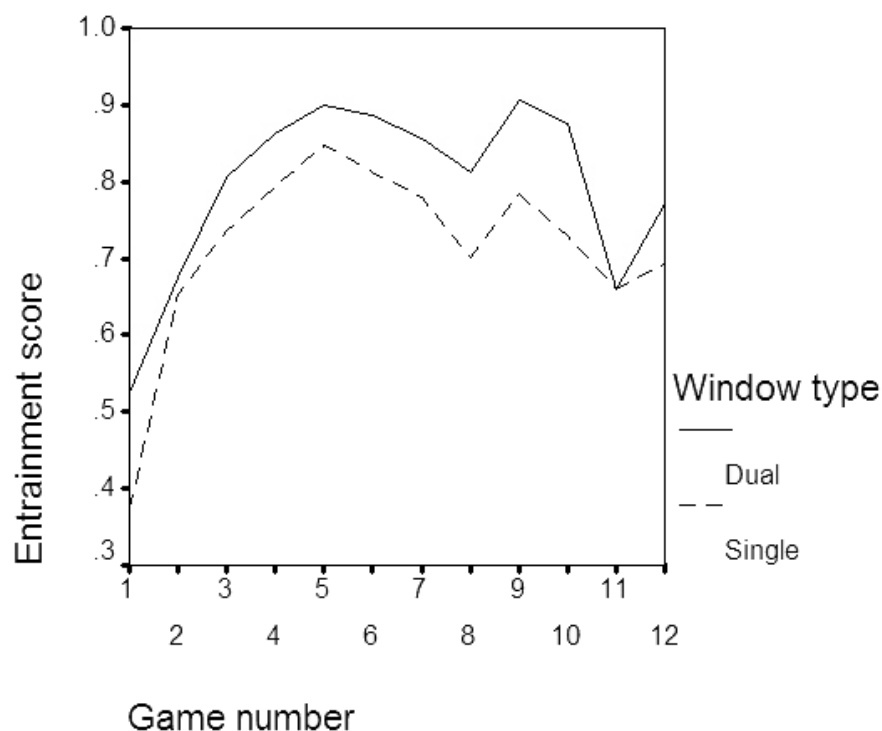


Figure 5.4: Entrainment scores of dual and single window dialogue

dual window turns showed a reliable main effect of Experience ($F_{(2,3747)} = 3.25, p = 0.04$) but no reliable reduction in turn-length (Linear trend, $p = 0.191$). The mean turn-lengths are presented in table 5.1.

	early (game 1-4)	middle (game 5-8)	late (game 9-12)
Single window	14.3	11	10.3
Dual window	14.6	13.4	13.9

Table 5.1: Length in characters of dual and single window turns

5.4.3 Effects of reduced sequential coherence

As a development of the initial question of how sequential coherence affects semantic co-ordination, four additional measures were used to test the hypothesis that participants should experience greater difficulty in exploiting the immediate local context.

One symptom of this greater difficulty in exploiting the local context should be that formulating questions should be more difficult in the dual window condition. Each turn was classified as a question or ‘normal’ chat-text. Manual classification was assisted by automatic detection of

participants' use of question-marks. An analysis of variance with Window type (dual/single) as fixed factor and Turn-type (question / non-question) as random factor showed no main effect of Window type ($F_{(1,8032)} = 1.72$, $p = 0.41$) and no main effect of turn type ($F_{(1,8032)} = 0.56$, $p = 0.59$) but a reliable two way interaction ($F_{(1,8032)} = 14.8$, $p = 0.00$). As figure 5.5 shows, questions are reliably longer in the dual window condition. As an extra check on the single window condition a further focused comparison on turn length (question / non-question) was carried out. A single-factor analysis of variance showed no reliable difference ($F_{(1,4284)} = 3.12$, $p = 0.08$).

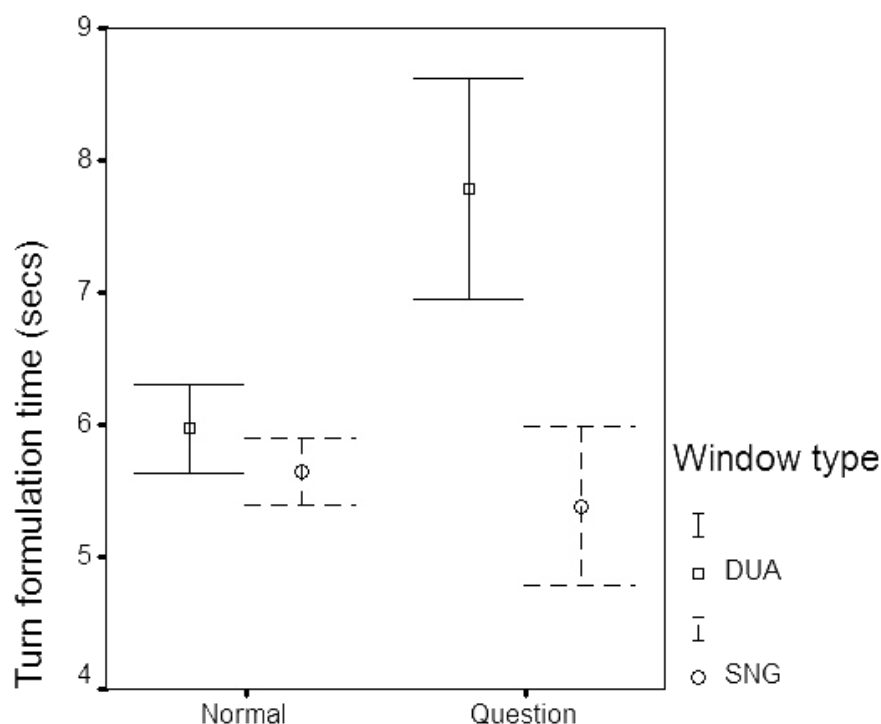


Figure 5.5: Turn formulation time of questions and non-questions in dual and single window dialogue.

Focusing on the relative frequency of one-word questions³ showed that there were proportionally fewer one-word questions in dual window dialogue ($Chi^2_{(1)} = 52.8$, $p < 0.01$). In dual window dialogue, 16 % of questions were single word, whereas in the single window condition 34% were one-word. By contrast, the proportions of one-word non-questions were broadly similar in both conditions (dual window 27% vs. single window 30%).

All turns were coded using the clarification request protocol of Purver et al. (2003). Multi-

³The category of one-word questions includes questions such as 'what?' as well as questions consisting only of question-marks, e.g. '???'.

nomial regression showed that there were fewer clarification requests in dual window dialogue ($Chi^2_{(1)} = 15.94$, $p = 0.00$). In contrast to Purver et al's and Schlangen's (2004) observations, where 5% of turns were found to involve clarification, in dual window dialogue the frequency was 1.5% and in single window dialogue 2.8%.

The frequencies of simple acknowledgements, e.g. 'right', 'ok' were also determined. Multinomial regression showed a reliable difference between dual window dialogue (4.3%) and single window dialogue (5.4%) ($Chi^2_{(1)} = 4.72$, $p = 0.03$).

5.4.4 Discussion

The results indicate that the use of the chat tool in combination with the maze game and the use of the modified chat tool to reduce sequential coherence in participants' dialogue was a qualified success. The global distribution of description types in the baseline condition replicates the patterns of use observed in spoken maze game studies (Garrod and Anderson, 1987; Garrod and Doherty, 1994). In particular, the pattern of migration from relatively concrete (Figural/Path) descriptions that depend on the specific details of each maze towards more abstract (Line/Matrix) description types that invoke schemata that generalise across instances is the same in both modalities.

Although participants take the same amount of time to complete the task in both conditions using a similar number of turns, the data show that the experimental manipulation has a pronounced effect on the pattern of description types in dialogue with reduced sequential coherence. Already in the first four games, participants in the baseline condition use more Matrix descriptions than in the dual window condition, a difference which increases in the last four games, in which over 75% of the location descriptions are of the Matrix type. This marked difference in the global distribution of description types between baseline and dual window conditions thus supports hypothesis 3. Of particular interest is that in the dual window condition participants initially use more abstract (Line/Matrix) descriptions than participants in the single window condition. This will be returned to below.

While in the baseline condition the global shift toward more abstract descriptions replicates the results of previous maze game experiments, when compared with the static pattern observed in dual window dialogue, tensions emerge in accounting for the differences between local and global alignment. Arguably, the static dual window pattern is evidence of less global alignment, since the most frequently used description type does not inhibit the use of other competing de-

scription types in a ‘winner takes all’ process (cf. Steels 2005). By contrast, in the baseline condition, participants rapidly adopt Matrix descriptions. Locally, this pattern of alignment is reversed, as baseline dialogue is associated with lower entrainment scores, confounding the basic predictions of the interactive alignment model that a decrease in sequential coherence should lead to less entrainment. Importantly, the entrainment scores are consistently higher in the dual window condition.

What could account for this marked difference? Before returning to the distribution of description types, the experimental manipulation is premised on the assumption that separating the chat-text into two separate windows should lead to participants encountering greater difficulty in establishing inter-speaker turn coherence. The data overwhelmingly indicate that this is indeed the case. While participants use less contracted turns in the dual window, these turns are typed faster than in the single window. A possible explanation is that participants in the dual window condition are freed from the constraint of producing interleaved turns, allowing them to formulate turns more quickly and in overlap, similar to the patterns observed in Anderson et al’s (2000) studies. These phenomena suggest that the overall effect of the reduction in sequential coherence is a shift away from fully interleaved dialogue towards dialogue that partially resembles a pair of ongoing monologues, disrupting both implicit and explicit attempts at turn-integration.

This appears to be reflected by dual window participants using fewer ‘Ok’ acknowledgements than in the baseline condition; these are highly context-dependent, since they do not specify what is being acknowledged (Allwood, 1992) and are consequently less effective if the immediate local context is not available. Closer analysis of the use of ‘Ok’ in dual window dialogue reveals the striking observation of the greater prevalence of questions (but not non-questions) that are prefixed with a turn-initial ‘Ok’, for example:

- (4.10) | (1) A: I have two gates also
 | (2) B: but I have only 1 switch after my gate
 | (3) A: *ok, are you stuck anywhere, so u need to let me through first?*

- (4.11) | (1) A: I’m on the third square from the top
 | (2) B: *ok wheres ur switch?*
 | (3) A: from there
 | (4) A: do
 | (5) A: right, right, down, left, left

It appears that this use of ‘Ok’ in questions is an adaptation to dialogue with reduced sequential coherence, since unlike spoken dialogue, where a question is usually adjacent to the

turn giving rise to that question, in the dual window condition, participants compensate for this inability to draw on the immediate local context by signalling with turn-initial ‘ok’ that the preceding dialogue has been understood, thereby creating a *structural boundary* which claims that the transition of this boundary is both expected and mutually agreed upon (Condon, 2001) (cf. Schegloff, 1982).

Participants’ compensation for not being able to rely on turn-adjacency is also evidenced by more explicit forms of demarcating boundaries in the dialogue. Excerpt 4.11 above highlights two important features: A’s three contiguous turns create a boundary between the two location descriptions, rendering the second location description more visible. This Path description, which is condensed into a single line, contrasts with the baseline and original maze game experiments where Path descriptions were generally used in installment sequences involving multiple acknowledgements by the other participant. This suggests that participants are orienting towards the lower effectiveness of their use, thus compensating for the increased attentional demands of identifying new turns and establishing relationships between turns in the separate windows.

A tentative explanation is that if participants are required to make the linkage between turns more explicit, and turns are generally less ‘noticeable’, new introductions of a description type will be less salient and consequently less likely to be adopted by the other participant, thereby resulting in a reduction of exemplars and hence in less stable representations.

Although this might explain the static global pattern in the dual window, this does not provide a satisfactory account of the higher entrainment scores, since these are consistently higher than in the baseline condition. This suggests that participants are not aligning more in order to reach convergence, but are aligning in response to the problem of reduced sequential coherence. This is the exact opposite of the interactive alignment model’s predictions that weaker co-ordination should result in less entrainment, since the lower frequency of clarification requests should, all things being equal, result in greater divergence of participants’ representations of the maze. As outlined in the introduction, turn-adjacency is integral to the successful deployment of clarification requests. Although clarification requests might be expected to signal greater communicative difficulties, they are less prevalent in dual window dialogue, suggesting that their use is inhibited by participants not being able to draw on the immediate local context. The data also show that formulating questions in general also requires greater effort. Importantly, if as the data suggests, participants are encountering greater difficulties in co-ordinating in the dual window condition,

this suggests that the problems encountered by participants in co-ordinating are exacerbated by this difficulty of formulating questions.

The question then arises as to why participants converge on Matrix description types, but rely on the Figural and Path description types when interaction is problematic (Healey, 1997). This presents a problem for both the interactive alignment model and the collaborative model, since the differences in description type are not reducible to the contraction of referring expressions; the Line and Matrix schemes do not emerge as abbreviated versions of Figural and Path schemes, rather they involve changes of semantic model.

In the CM and the IM, the development of co-ordination is characterized as initially involving solutions that are often based on salience. If successful, these initial attempts can form the basis for co-ordination on precedence, of which the repeated use allows the development of more stable conventions that are resistant to individual misapplication. Co-ordination on precedence in the interactive alignment model depends on the availability of the immediate local context for the priming of interlocutors' production and comprehension mechanisms. This assumption is not confirmed by the data, as the very manipulation that should interfere with this process of local input-output matching has the opposite effect of increasing the entrainment between participants' location descriptions.

This clearly suggests that local matching *per se* cannot be the main mechanism underlying global alignment since, in the baseline condition, the entrainment scores are consistently lower. Similarly, the global dual window pattern cannot be explained by participants' attempts at developing a convention being confounded by individual violations which lead to the nascent convention losing its justification for its use, since this should also be reflected in lower entrainment scores.

There is also the additional problem of accounting for dual window participants' initial use of the more abstract Line and Matrix descriptions since the reduced sequential coherence should result in participants experiencing greater difficulty in providing positive and negative evidence of understanding, and on this basis should be more weakly co-ordinated. However, the pattern of description types used by dual window participants indicates that they are, in effect, initially treating each other as more highly co-ordinated than single window participants.

Taken together with the post-hoc tests, these apparent contradictions between participants' co-ordination on salience and precedence suggest that they are secondary effects of the manipula-

tion of sequential coherence and provide convergent evidence that questions, in particular those that seek to clarify misunderstandings concerning the use of maze game descriptions, are the principal locus of the effect of sequential coherence on semantic co-ordination. The importance of questions is underscored by their increased turn-formulation times in dual window dialogue; questions are *prima facie* directly tied to the immediate local context, whereas acknowledgements are less sensitive to what is being acknowledged.

Bringing the signalling of misunderstandings into the foreground allows an account of dual window participants' initial use of more abstract descriptions. By hypothesis, this difference is attributable to participants being less able to signal the extent of their (mis)understandings concerning the use of abstract descriptions: in both conditions, participants initially attempt similar proportions of concrete and abstract descriptions. In the baseline condition, participants are able to clarify the problems they encounter in using abstract descriptions, either resolving the misunderstandings and continuing their use of abstract descriptions, or resorting to less abstract descriptions that depend on the concrete details of the maze layout, hence allowing simpler co-ordination through salience. In the dual window condition, however, the 'slippage' introduced by the reduction in sequential coherence interferes with the ability of participants to both detect problems that arise in using the abstract descriptions, but most importantly, increases the effort required to signal problems and ask each other questions, since participants cannot draw on the immediate local context, resulting in ambiguities and misunderstandings concerning the use of abstract description types that persist for longer in the dialogue. Inspecting the transcripts yielded two dyads in the dual window condition that clearly demonstrate this persistence of misunderstandings concerning the use of abstract descriptions. The first dyad used the Matrix scheme for the first 3 games, despite both members counting from opposite corners of the maze. The second dyad, although counting from the same origin, used a combination of Line and Matrix descriptions for 5 games before establishing that one member had consistently been counting from 0, whilst the other had been counting from 1.

The data demonstrates a strong effect of the interleaving of chat-text on the development of co-ordination, revealing interactional contingencies between the kinds of description used in the maze game and the possibility of engaging in clarification. The data suggest that it is not simply the provision of feedback that is implicated in this process, since the chat tool does not block any of the turns and participants are still able to compensate through providing explicit

forms of positive and negative evidence of understanding. Instead, it appears that participants' turn-by-turn modification of each other's turns is critical to the development of co-ordination.

5.5 Conclusions

Two important questions arise. Firstly, why does dialogue that is beset by communicative difficulties result in more entrainment? The data suggest a more complex view of the role of alignment: in the IM it is seen as a fundamental precondition for co-ordination, yet here it appears that participants are aligning locally in response to communicative difficulties, which indicates that one of the techniques that participants use when they are not able to formulate questions is to compensate by means of alignment.

The second question concerns the frequency of occurrence of clarification requests in the dialogue. The results strongly suggest that participants' opportunities for engaging in repair is fundamental to the development of semantic co-ordination, yet there are far fewer clarification requests in the baseline condition than in previous maze game studies. Although this paucity can in part be explained by the persistence of the chat-text resulting in fewer clarification requests that address attentional or lexical problems, in analyses of spoken corpora 5% of utterances were found to involve clarification requests (Purver et al., 2003; Schlangen, 2004). This will be returned to in chapter 7.

Chapter 6

The effect of participatory status on co-ordination

6.1 Introduction

The results from the previous experiment emphasize the importance of direct interaction in the development of co-ordination. Although the results indicate that participants' opportunities for engaging in clarification play a key role in this process, the experimental manipulation in the dual window not only interfered with participants' ability to clarify, but also interfered with interspeaker turn coherence. This manipulation raises the possibility that some of the problems encountered by dual window participants are attributable to this artificial decrease in turn-relevance and adjacency. There is thus the concern that the increased need for clarification created by introducing additional communicative demands are less representative of the co-ordination problems faced in 'ordinary' maze game dialogue.

What is required, therefore, is to observe the effects of different levels of direct interaction without introducing systematic, continuous disruption of the turn-by-turn provision of feedback between participants, and instead to manipulate more directly participants' opportunities for clarification. A proven method in previous studies is to assign participants different levels of active participation, thereby manipulating their opportunities for direct interaction. This approach was adopted in the current experiment by using the chat server to generate artificial questions about location descriptions which appeared to originate from participants of different levels of participation. Before describing the experiment, the discussion of the experiments conducted by Clark et al. (section 3.5.1) will be resumed to assess their implications for the development of

co-ordination in the maze game.

6.2 Participatory status

In the collaborative model differences in participatory status are analysed as differences in the obligations and opportunities that interlocutors have to engage in the grounding cycle. When conversing in the presence of side-participants, bystanders and eavesdroppers (peripheral participants), the collaborative model suggests that the ratified participants of the conversation (primary participants) draw on their knowledge of the common ground that exists between all participants, pursuing a strategy of either indifference, disclosure, concealment or disguise (Clark and Schaefer, 1992). This is characterized as a complex process of calculations and exploitations of the common ground, although it is less complex if a strategy of indifference is adopted.

Experimental evidence for differences in co-ordination developed by interlocutors of different participatory status is provided by the experiments conducted by Schober and Clark (1989), in which peripheral participants, although exposed to the same informational content as addressees, were consistently slower in their development of co-ordination. The collaborative model explains these findings by peripheral participants being less able to co-ordinate effectively as they are restricted in providing positive and negative feedback to the primary participants, which limits their involvement in the contribution cycles of the dialogue. The CM also draws on experimental evidence indicating that participants actively index each other's participatory status and formulate their messages accordingly. Wilkes-Gibbs and Clark (1992) propose that even prior to directly interacting with peripheral participants who have overheard the preceding conversation, speakers' first utterances already reflect assumptions about peripheral participants' level of grounding and their associated levels of copresence.

By contrast, the interactive alignment model presupposes that interlocutors do not actively index each other's level of understanding, relying instead on the overlap between representations that is provided by convergence due to priming. Although some concessions are made to allow the 'postconscious' (Bargh, 1989) influence of participatory status, all things being equal, it denies that interlocutors dynamically accommodate each other's utterances with inferences based on full common ground and thus encounters difficulties explaining the differences between primary and peripheral participants observed in the experiments described above. However, Carletta, Garrod and Fraser-Krauss (1998) argue that the priming mechanism can explain

differences in multi-party interactions: peripheral participants have fewer opportunities to address problems they have in comprehending a primary participant's contribution, resulting in weaker co-ordination due to the peripheral participants being exposed to fewer instances of each input.

In summary, both the CM and the IM attribute differences in semantic co-ordination achieved in multi-party exchanges to differences in participants' opportunities for interaction. Both predict that active participants will co-ordinate more strongly and more quickly than peripheral participants. However, they differ in their predictions about how interlocutors will respond to these differences. The CM predicts that speakers actively track the different levels of co-ordination that develop with different participants. The interactive alignment model predicts that speakers respond instead to the cumulative exposure to particular inputs independently of their origin in the conversation.

To test these predictions the maze game and chat tool described in section 4.3.3 were adapted to introduce artificial probe questions about spatial locations within the maze game. The questions were generated to appear to participants to originate either from each other or from a third peripheral participant.

6.3 Methods

The chat tool and the maze game are the same as described in section 4.3.

6.3.1 Chat server

The chat server generates artificial probe questions which appear, to participants, to originate either from each other (primary probe) or from an overhearer (peripheral probe). Each turn is preceded by the name of the apparent source (either the other participant's chosen nickname or the experimenter's name) followed by a colon. Six interventions were used, each designed to elicit a spatial description and to be plausible in a range of dialogue contexts:

1. Where's your goal?
2. Where's the gate nearest to your goal?
3. Where's your nearest gate?
4. Where's the switch nearest to your goal?

5. Where's your nearest switch?
6. Where did you start from?

Probes are sent simultaneously to both participants. This helps to minimise disruption to the dialogue by keeping both participants engaged in interaction and ensuring that both experience similar patterns of questioning during the task. Probes are dynamically modified to mimic the speed of typing and spelling / 'txt' conventions used by each participant.

The responses to the probes are captured by the server. The probe and the participants' response are displayed only in their own chat-window, the other participant does not see them. In order to co-ordinate the resumption of the interaction after a probe, the server monitors whether one participant starts typing before their partner has finished responding to the probe. If this occurs, an error message is displayed and further text-entry is prevented until either the partner responds or a pre-defined time-out expires. Subsequent turns are relayed as normal. To ensure error messages do not cue the interventions a small number of random error messages are also introduced at other points.

6.3.2 Participants

31 pairs of participants were recruited, 38 male and 24 female, from undergraduate students. They were recruited in pairs to ensure they were familiar with each other. Only participants who had some previous experience of using internet chat software such as ICQ or Microsoft Messenger were selected for the experiment. Each participant was paid £10.00 for participating in the experiment.

6.3.3 Procedure

The same basic procedure was followed as outlined in section 4.3. To ensure participants had some experience of responding to queries apparently originating from the experimenter (peripheral participant) an initial probe asked participants if they could read the text. On receipt of an acknowledgement, a turn instructing participants to start the experiment was sent. Twelve mazes were presented in random order to each pair. Probe questions were introduced only in the first four and the last four games. This created two conditions: 'early' and 'late' which indexed different levels of exposure to the dialogue. Half of the probes had the other (primary) participant as the apparent source and half the experimenter (peripheral) as the apparent source. 32 randomly

selected probe questions were used with a maximum of two interventions per participant per maze. In addition, no intervention from the same apparent origin was repeated in the first four or last four games. Overall, this resulted in a factorial design with two within-subjects factors: Source (primary vs. peripheral) crossed with Exposure (early vs. late).

6.4 Results

On debriefing, no participant reported detecting the artificial probe turns. Times from the log files provided two measures of response to the interventions. Firstly, turn formulation time, the time from the onset of typing of a response to a probe turn to its completion was calculated. This was analysed in a 2x2 analysis of variance with Exposure (early vs. late) and Source (primary vs. peripheral) as within-subjects factors. There was a main effect of Exposure ($F_{(1,412)} = 8.61$, $p = 0.04$), no main effects of Source ($F_{(1,412)} = 0.09$, $p=0.77$) and no interaction ($F_{(1,412)} = 0.25$, $p = 0.62$). Overall, participants became faster at producing their responses over time, taking an average of 23 seconds in the early trials and 18 seconds in the late trials.

The second measure of task performance used was latency of response to the probe turns: the time between the onset of an intervention and the initial onset of typing the response (regardless of whether there was subsequent deletion) was extracted from the logs. A 2x2 analysis of variance with Exposure and Source as within-subjects factors showed no effect of Exposure ($F_{(1,412)} = 0.08$, $p = 0.77$), a main effect of Source ($F_{(1,412)} = 10.56$, $p=0.001$) and no interaction. Initiation of responses to peripheral probes took twice as long (12 secs) as responses to primary probes (5.9 secs) and this difference was consistent throughout.

6.4.1 Description types

For comparison with previous work a total of 9755 turns, including both descriptions generated in normal dialogue and 684 responses to probe turns were classified according to the criteria described in section 5.4.1.

Transcription results

The global change in the distribution of description types in spontaneous dialogue (excluding responses to probe questions) is illustrated in Figure 6.1.

The distribution of description types used in responses to interventions is illustrated in Figure 6.2. This suggests a different pattern of responses to primary and peripheral participants. Multi-




Figure 6.1: Distribution of description types in spontaneous dialogue in the first four games (early) and the last four games (late).

nomial regression shows a significant main effect of both Exposure ($Chi^2_{(3)} = 14.5$, $p = 0.00$) and Source ($Chi^2_{(3)} = 15.6$, $p = 0.00$) on choice of description types.

To provide focused tests of the hypotheses three additional comparisons were carried out using multinomial regression. Firstly, the prediction that peripheral participants' level of coordination should change over time was tested by comparing the distribution of description types produced for peripheral participants only in early vs. late trials. This showed no reliable difference ($Chi^2_{(3)} = 0.82$, $p = 0.66$). Secondly, a test of whether the description types produced for primary and peripheral participants differ in the early games showed a reliable effect of Source ($Chi^2_{(3)} = 15.5$, $p = 0.00$). Finally, a test for whether primary and peripheral differ in the Late games showed no reliable difference ($Chi^2_{(3)} = 3.98$, $p = 0.26$). Overall, the profile of description types produced for primary participants evolves over trials whereas those produced for peripheral do not change. This pattern is illustrated in Figure 6.2.

6.5 Discussion

The distributions of description types observed here replicate the basic patterns observed in the original, oral, maze game studies and the previous experiment. There is a general migration from the relatively concrete descriptions (Figural and Path) that depend on the specific details of each maze, towards more abstract description types (Line and Matrix) that invoke schemata that generalise across instances. This is true both of the descriptions produced in the spontaneous dialogue (Figure 6.1) and for those produced in response to the experimental interventions with

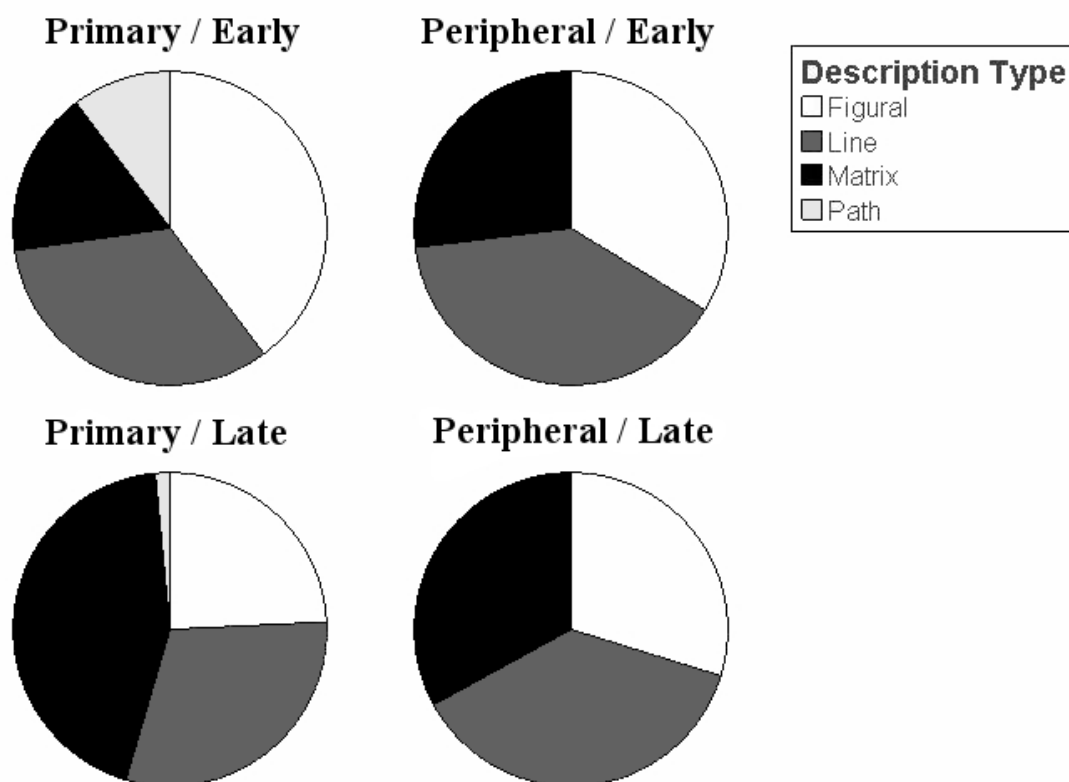


Figure 6.2: Distribution of description types in response to artificial questions by primary and peripheral participants

the primary participant as their apparent source 6.2.

As noted in the introduction, Carletta et al. (1998) propose that the same basic mechanisms of alignment, i.e. input-output co-ordination, should operate in dyadic and in multi-party interaction, modulo opportunities to interact. Participants should use the lexical, semantic and syntactic forms which are most strongly primed by the preceding dialogue.

The results reported here show, however, that maze game participants reliably distinguish between primary and peripheral participants. The same questions receive different responses depending on whether the apparent source of the question is a primary or peripheral participant. Specifically, they take twice as long to initiate responses to probes that appear to originate from a peripheral participant and, in the early exchanges, they use different description types. As a result the local choice of description type in the responses to probes is not explained only by reference to priming from the preceding dialogue.

The collaborative model appears to be better equipped to deal with these observations since it assumes that levels (or forms) of co-ordination will be explicitly indexed to different partici-

pants (Brennan and Clark, 1996). However, all things being equal, levels of co-ordination should be lower where participants have not explicitly engaged in a cycle of grounding (Clark, 1996b; Wilkes-Gibbs and Clark, 1992). In the experiment reported here the peripheral participants provide no positive evidence of acceptance and on this basis should be more weakly co-ordinated. However, the pattern of description types used indicates that peripheral participants are actually treated as more highly co-ordinated in the early games than primary participants.

One possible explanation for this is a ‘lab coat’ effect. Participants might assume that the peripheral participant, the experimenter, is a maze game expert who can understand more complex forms of description (Clark, 1996b). If participants have a conception of what constitutes ‘expert’ language they could use it with the experimenter from the start of the task. However, it then becomes difficult to explain why they do not do the same thing with their task partner since, by hypothesis, they can both already understand and produce the more ‘expert’ descriptions and, unlike the experimenter, are actually able to ground them with their partner.

The patterns observed here and in the previous experiment provide convergent evidence of the key role of the signalling and resolution of misunderstandings in the development of semantic co-ordination. Specifically, the initial treatment of peripheral participants as more co-ordinated is analogous to the initial use of more abstract descriptions in the dual window condition of the first experiment. Building on the explanation of the role of negative feedback in co-ordination, the initial difference between the primary and peripheral participants appears to be due to primary participants being able to provide feedback to each other about problems they are encountering. In response they switch to more basic description types and over time build up their co-ordination. By contrast, peripheral participants provide no feedback about the use of the description types, and hence provide no impetus for switching to the more basic concrete descriptions in order to resolve the difficulties.

In addition to these local patterns, the maze game data also display more global patterns of co-ordination. As observed in the previous experiment and in the original maze game studies, participants migrate from the more concrete instance-specific description types to more abstract ones that capture generalizations about the underlying grid structure of the maze. These more global trends also pose a problem for both models. They are not accounted for by priming since, as Figures 6.1. and 6.2. illustrate, they run counter to local precedence. Over time, participants do not consistently converge on the description type that they are most frequently exposed to:

here, as in the sub-community group described in (1994), participants initially use more Line descriptions before converging in the later games on Matrix descriptions. As mentioned in section 3.7 these trends have a different character to the patterns of contraction or abbreviation of referring expressions that are the primary concern of the collaborative model.

Relegating the importance of alignment as the pre-condition for co-ordination, and bringing the mechanism of repair into the foreground allows a different conception of how semantic co-ordination develops. By hypothesis, co-ordination is built on the opportunities interaction creates for identifying and addressing *differences* in interpretation. Initially participants try out whatever description type occurs to them. Once they detect a problem with interpretation they systematically exploit less abstract schemes as one means of repairing the problem. This allows them to take advantage of the local context, in particular the maze in front of them, to resolve uncertainties about what location was intended. By repeatedly resolving ambiguities in this way participant's semantic models can progressively converge. This is, in effect a semantic variant of Clark and Marshall's (1981) concept of vertical repair. However, the evidence suggests that the basic strategy for dealing with communication problems is to switch to less, not more, specific forms of description.

6.6 Conclusions

This experiment demonstrates that there are differences between primary and peripheral participants in terms of the speaker's choice of description type. This effect appears to depend, as in the previous experiment, on the kind of direct interaction that is possible. In this experiment, the level of inter-turn coherence of primary participants' dialogue was not directly manipulated, but still resulted in effects that are also apparently to do with participants' failure to go through a process of genuinely determining whether co-ordination has actually been achieved. What might the locus of these effects be? The account developed so far strongly suggests that participants' opportunities for directly resolving misunderstandings are of key importance to this process.

Chapter 7

The role of clarification requests in semantic co-ordination

7.1 Introduction

The two previous experiments both point towards the importance of repair in the development of semantic co-ordination. The results indicate that one of the ways that participants introduce abstract descriptions is by attempting their use without being sufficiently co-ordinated to use them *ab initio*, requiring the use of more concrete, instance-bound descriptions that allow ambiguities and confusion surrounding their use to be resolved. It is strongly suspected that this is achieved via the use of repair, which depends on participants being able to draw on the immediate local context, exploiting the sequential coherence of the dialogue in order to signal and resolve difficulties.

While the evidence is compelling, it was obtained indirectly and the manipulations also interfered with other co-ordination mechanisms, in particular with participants' provision of positive evidence of understanding. Thus it is unclear how clarification requests are deployed in order to resolve the different kinds of misunderstanding that arise and how the global shift towards abstraction is reflected within these clarification sequences.

7.2 Clarification requests

A common thread running through the clarification taxonomies proposed by Schlangen (2004), Gabsdil (2003) and Purver (2003) is that different clarification requests access different levels of

understanding within an action hierarchy as described in section 3.5.2

- (7.1) | (1) A: I'm on the second switch
 | (2) B: switch?
 | (3) A: yeah the grey thing

So, for example, in example 7.1 above, the clarification request 'switch?' might typically signal problems in recognizing the meaning of 'switch', requesting further specification, whereas a clarification such as 'what?' is generally associated with misunderstandings occurring at the lower levels of securing attention or recognition of the utterance. The choice of clarification type can thus signal the information required for understanding to occur at all levels of the hierarchy. Clarification requests (henceforth CRs) such as 7.1 above typically signal more co-ordination than 'open-class' (Drew, 1997) repairs such as 'what?', as the former give more clues about the nature of the problem and consequently the expected response. By contrast, 'open-class' repairs provide no specific information about the nature of the problem being signalled, and are even more reliant on adjacency to the problematic turn for their successful deployment.

The ordering in strength of CR types trades on the pragmatic expectation that interlocutors, in seeking to minimize the 'collaborative effort' will normally design their CRs to give as much information as possible about their current level of understanding. Although 'what?' can be used to clarify at higher levels (Drew, 1997), the expectation is that interlocutors should produce CRs that signal the highest level of understanding currently available to them, which in the collaborative model is formulated as the 'strongest initiator rule' (see section 3.5.2).

Hierarchical approaches can thus categorize and rank sources of problematic understanding, however their 'semantic neutrality' precludes their provision of an account that is sensitive to semantic differences between description types. When participants change schemes, the interactive alignment prediction is that the most frequently used (i.e. primed) scheme will predominate. This conservative notion of repair would appear to stand in contrast to the observed shift toward more abstract descriptions, which in previous maze game experiments and in the second experiment ran counter to local precedence.

Similarly, there is a general expectation in both Purver et al's (2003) and Rodriguez and Schlangen's (2004) models that interlocutors will modify their original utterance in response to a CR. However, where the reformulation occurs within a particular level of the hierarchy, there are no mechanisms for predicting what kinds of semantic change will occur in response.

In summary, the previous experiments strongly suggest that the opportunities for direct interaction contribute strongly to the development of semantic co-ordination, in particular the opportunities for using clarification requests in resolving misunderstandings. However, it is unclear how participants' choice of CR is motivated by the particular kind of misunderstanding and how the local responses to the CRs enable the global shift toward more abstract descriptions.

The experiment was designed to address the basic empirical questions whether there is a direct connection between the occurrence of CRs and semantic co-ordination and whether there is a connection between the 'level' of CR and the form of semantic co-ordination. The basic rationale of the experiment was to test the effects of different CR types on the form and content of participants' responses. Before introducing the specific hypotheses, the experimental methods are explained in more detail below.

7.3 Methods

Both the chat tool and the maze game are the same as described in section 4.3.

7.3.1 Chat server

In addition to relaying turns between participants, the server monitors the content of the turns in order to generate artificial clarification requests that appear, to participants, to originate from each other.

The server compares each turn with a lookup table of location descriptions obtained from 10000 turns from the previous experiments, combined with rules for detecting misspellings and non-standard 'txt' conventions. This ensures that CRs are generated only on turns containing spatial descriptions.

Each clarification request generated by the server is preceded by the other participant's chosen nickname, followed by a colon, and is dynamically modified to mimic spelling and typing speed.

To provide a manipulation of CR type, the following two classes of CR were selected: Reprise fragments ('Fragments') that echo a word from the target turn and 'Whats' (e.g. 'what?' or 'sorry?') that query the turn as a whole. These are the two most common forms of CR in ordinary dialogue (Purver et al., 2003) and they provide two different levels of clarification. Reprise fragments involve direct re-use of a word from the turn and imply that the rest of the turn was understood.

By contrast ‘Whats’ suggest that there were global problems finding a sense for the turn (but see also Drew (1997)).

‘Frag’s (High co-ordination): Repetition of a single fragment of the location description.

‘Whats’ (Low co-ordination): What? Huh? Sorry? Ehh? Where? Sorry what?

Participants’ responses to the probe CRs are captured by the server. The probe CR and the response are displayed only in the participant’s own chat-window. After receiving a response to the CR, the server sends one of the following acknowledgement turns to the recipient: ‘ok’, ‘k’, ‘ok right’ and resumes relaying subsequent turns as normal. During the artificial CR exchange the server monitors whether the other participant starts typing. If this occurs, an error message is displayed and further text-entry is prevented until either the CR sequence is finished or a pre-defined time-out threshold is reached. To ensure error messages do not cue the interventions, a small number of random error messages are also introduced at other points in the dialogue.

7.3.2 Participants

21 pairs of native English speaking participants were recruited, 23 male and 19 female, from undergraduate students. They were recruited in pairs to ensure that they were familiar with each other. Only participants who had some previous experience of using internet chat software such as ICQ or Microsoft Messenger were selected for the experiment. Each participant was paid £10.00 for participating in the experiment.

7.3.3 Procedure

The same basic procedure was followed as described in section 4.3.4.

Twelve mazes were presented in random order to each pair. Artificial clarification requests on turns that involved spatial descriptions were introduced throughout the experiment with a 5% probability of any turn being clarified, consistent with findings from Purver et al. (2003) and Schlangen (2004) that approximately 5% of dialogue turns are CRs.

The experimental group were thus exposed to a within-subjects manipulation of CR type (‘Frag’ vs. ‘What’). The ‘single window’ group from the first experiment was used as a baseline control group.

7.3.4 Experimental hypotheses

1. The introduction of artificial CRs will interfere with semantic co-ordination.
2. More severe problems will cause more disruption (i.e. 'Whats' will cause more disruption than 'Fragments')
3. Participants will systematically shift to more concrete forms of semantic co-ordination (Figural / Path descriptions) where problems occur.

7.4 Results

Overall, 246 clarification requests were artificially generated by the server: 109 'Fragments', 128 'Whats' and 9 CRs generated for turns that did not contain spatial descriptions. These non-spatial clarifications were excluded from further analysis. On debriefing, no participants reported detecting that the probe CRs did not originate from their partner.

7.4.1 Description types

Both the target turns used by the server to generate CRs and participants' responses were classified according to the criteria described in section 5.4.1.

Figure 7.1 illustrates the contrast in the global distribution of description types in the baseline control condition and in dialogue that is periodically interrupted with artificial clarification requests. The difference is reliable (Multinomial Regression: $(Chi^2_{(3)} = 276, p = 0.00)$). The results show that the probe CRs significantly disrupt co-ordination in the experimental group. The largest category of description type in the experimental group is Figural whereas in the baseline control group the Matrix descriptions predominate.

To check whether co-ordination was still developing over time (but to a lower level) in the experimental group the distribution of description types used in target turns used in the first four games was compared with those used in the last four games¹ (see Figure 7.2). This showed that there was still a significant shift in the use of description types over time (Multinomial Regression: $(Chi^2_{(3)} = 15.1, p=0.00)$) with participants migrating from concrete descriptions (Figural/Path) towards abstract (Line/Matrix) in the later games. This suggests that semantic co-ordination was still developing but at a significantly slower rate than in the control group.

¹Target turns only were selected for this analysis as they are furthest from the immediate influence of the artificial CRs.

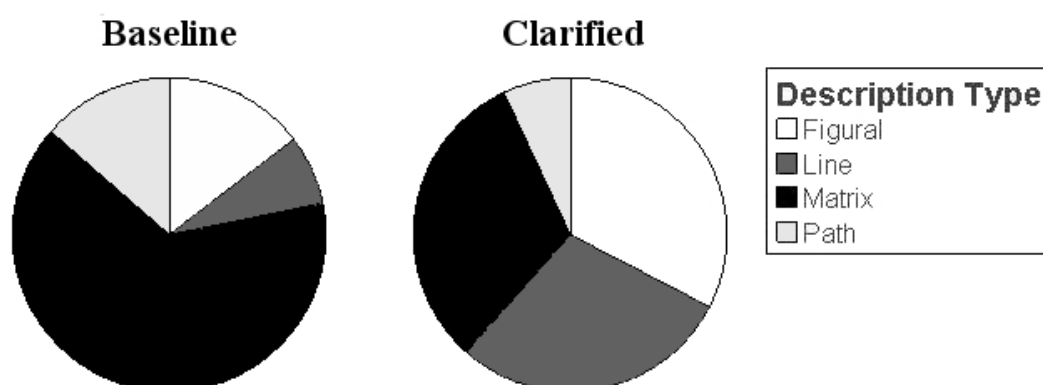


Figure 7.1: Distribution of description types in baseline and clarified dialogue

In order to test the effects of the two CR types on the way responses were formulated, a focused comparison of the distribution of description types in the responses immediately following the ‘What’ vs. ‘Frag’ CRs was made. This showed no reliable difference (Multinomial Regression ($Chi^2_{(3)} = 1.68, p=0.64$))




Figure 7.2: Distribution of description types of target turns in the first four games (early) and the last four games (late)

To provide an additional test of the third hypothesis, namely that participants systematically shift to more concrete description types as a way of resolving co-ordination problems, the relationship between the description type used in the target turn produced by a speaker and the spatial description type they produced in their response to the probe CR was examined. Out of a total of 142 spatial description responses 101 (71%) responses used the same description type as the target. Of the 29% (41) that changed description type 14 (34%) involved a shift from concrete




Figure 7.3: Innovation scores of all turns excluding CR responses

(Figural/Path) to abstract (Line/Matrix) whereas 27 (66%) involved a shift in the opposite direction. Overall, responses to the CRs predominantly used the same description type but where a change occurred it was more likely to involve a change to a more concrete description type.

7.4.2 Redundancy and innovation

A post-hoc analysis of the constituent words of location descriptions was carried out to provide measures of response innovation and redundancy. The ‘open-class’ constituent words of each location description (determined with WordNet) were used to calculate the number of turns separating each successive use of these constituent words. For example, a description such as ‘the middle line’ leads to a search backwards in the dialogue for the previous occurrence of ‘middle’ and ‘line’, recording the turn distances for both. Figure 7.3 shows a decrease of innovation over time, which fits the basic predictions of both the CM and the IM, with participants relying more and more on previously grounded or entrained expressions.

To take into account the increasing significance of newly introduced words as task experience progresses, the proportion of newly introduced words per location description was multiplied by the turn number of the location description. Thus if in the 20th turn there were 5 words, four

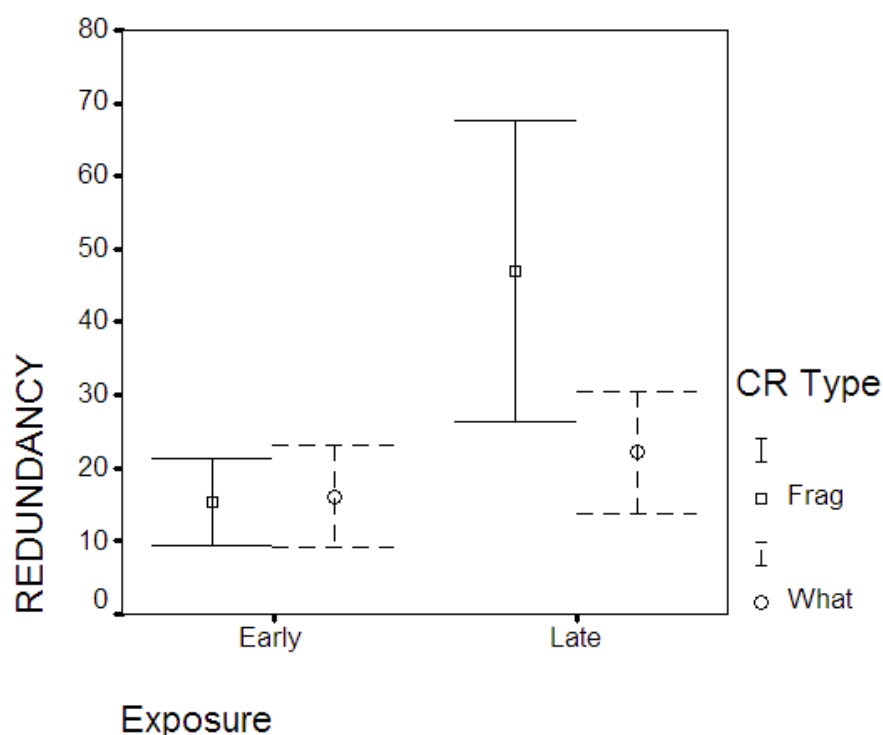


Figure 7.4: Redundancy of responses to 'Whats' and 'Frag'

of which were novel, the innovation score would be 16. Focusing on the differences between 'Whats' and 'Frag', a two-way analysis of variance yielded a significant effect of CR type ($F_{(1,200)} = 5.1$, $p = 0.025$), no effect of Experience ($F_{(1,200)} = 0.35$, $p = 0.56$) and no interaction ($F_{(1,200)} = 0.193$, $p = 0.61$). Overall, 'Whats' resulted in more innovation than 'Frag' throughout the experiment.

A third measure focused on the 'redundancy' of CR responses. This measure took all words in the response that had been previously used in the dialogue and calculated the average turn-distance separating their previous use. The development over time of the redundancy scores of both types of CR response was analysed in a 2x2 analysis of variance, yielding no main effect of Exposure ($F_{(1,200)} = 0.87$, $p = 0.52$), no main effect of CR type ($F_{(1,200)} = 2.16$, $p = 0.38$), but a reliable interaction ($F_{(1,200)} = 6.56$, $p = 0.01$). As shown in Figure 7.4, although initially similar, towards the end of the experiment the words reused in responses to 'Frag' are repeated from further back in the dialogue history than responses to 'Whats', which draw on turns that are closer to the CR.

7.4.3 Other measures of CR effects

Times from the log files were used to provide two further comparisons of responses to the ‘Frag’ and ‘What’ CRs. Firstly, turn formulation time, the time from the onset of typing of a response to its completion. A one-way analysis of variance revealed reliable differences between CR types ($F_{(1,235)} = 6.3$, $p = 0.01$). Overall, participants took longer to formulate their responses to ‘Whats’ than to ‘Frag’, taking an average of 25 seconds to respond to the former, and 18 seconds to respond to the latter.

The second measure of response time used was typing-onset time: the time between the onset of an intervention and the initial onset of typing the response. A one-way analysis of variance showed no effect of CR type ($F_{(1,235)} = 0.32$, $p=0.57$).

In order to provide a measure of the indirect disruption caused by a CR, data from the log files was used to calculate the number of turns between receipt of a CR by a participant and the next turn in which they produced a spatial description. A one-way analysis of variance revealed significant differences between CR types ($F_{(1,93)} = 8.46$, $p = 0.02$). Overall, ‘Frag’ caused less disruption (average 3.2 turns before next description) than ‘Whats’ (average of 5.3 turns before next description).

The log files were also analysed for number of ‘deletes’ or edits that occurred in the construction of a turn prior to sending it. Although there was no reliable evidence of a relationship between edits and description type in responses ($Chi^2_{(1)} = 0.881$, $p = 0.35$) there was a reliable relationship between edits of the target turn and description type of the subsequent response to the CR ($Chi^2_{(1)} = 9.9$, $p=0.002$). If there were no edits in the target turn there were more abstract responses (44 Matrix / Line vs. 23 Figural/Path). If the target was edited prior to sending, there were fewer abstract responses (32 Matrix/Line vs. 43 Figural/Path).

7.5 Discussion

The global distribution of description types reported here replicates the patterns of use observed in the two previous experiments and the original maze game experiments, as evidenced by the pattern of migration from concrete to more abstract description types.

While pairs in the baseline group from the first experiment converged on the ‘Matrix’ scheme, those exposed to CRs do not. Although their form of co-ordination does change over time it evolves more slowly and they do not converge on the Matrix scheme by the end of the experiment.

However, as in the second experiment, they do not converge on the initially most frequently used description types. The marked difference in the distribution of description types between the experimental (CR) and control (baseline) groups thus supports hypothesis 1.

The second question raised in the introduction was whether there is a connection between the particular type or ‘level’ of CR and the form of semantic co-ordination. The results reported here do not provide a clear answer to this question. Hypothesis 2 predicted that ‘Whats’ would cause more disruption than ‘Frag’s’; however no reliable difference was found in the distribution of description types in responses to the two CR types. This is not, however, because the two CR types failed to have any distinct effects. The response time data show participants took longer to formulate their responses to ‘Whats’ than to ‘Frag’s’, and the ‘disruption’ data indicate that participants took longer to get the dialogue back on track after a ‘What’ than a ‘Frag’. More importantly, the measures of redundancy and innovation show clear differences between the two CR types. This will be returned to below.

Overall, participants were sensitive to the difference between the two classes of CR. As expected, the ‘Whats’ were more disruptive to the dialogue than the ‘Frag’s’. However, while the evidence thus supports hypothesis 2, the results make it more difficult to explain the nature of the connection between CRs and forms of semantic co-ordination.

The third hypothesis considered above was that the local effect of CRs should be to prompt a shift from abstract (Matrix/Line) to concrete (Figural/Path) descriptions. The results provide some support for this. In the cases where participants do change description type in response to a CR, there is a greater preference for changing from abstract to concrete than vice versa. However, the more striking observation is that in 71% of cases participants do not change description type. This local consistency in description type echoes Garrod’s original findings, suggesting that the repairs are occurring ‘within’ the same situational model as the target-turn. However, it presents a puzzling contrast with the global effects of the CRs. Although the additional clarification questions have a significant impact on overall co-ordination, as indicated by choice of description type, it appears that these effects are not manifest in the choice of description type in the immediate context in which the CRs occur.

This local / global contrast is complicated by the patterns of redundancy and innovation that occur at the sub-description type level. Whilst ‘Whats’ consistently result in more innovation than ‘Frag’s’, they differ in their reuse of items from the dialogue history: once participants be-

come more co-ordinated, their responses to ‘Fragments’ reuse items from further back in the dialogue than their responses to ‘Whats’, which reuse items from the more immediate local context. Prima facie, this is an inversion of the IM’s predictions that the less severe misunderstandings should be resolved locally using the implicit common ground, since the responses to ‘Whats’ are more redundant than responses to ‘Fragments’.

It is unclear exactly how to characterize these different patterns of reuse, as ‘Whats’ appear to be both more innovative and redundant, suggesting that they engender at least two very different kinds of response. One possible explanation is that the ‘Whats’, in foregoing explicit diagnosis of the kind of problem, are essentially signalling that appealing to the local context is of no use in resolving the problem. On this view, although the ‘Fragments’ provide a diagnosis, as the participants become more co-ordinated, the diagnosis becomes less and less interpretable within the local context, as the terms within the local context seem to be mutually-intelligible and functioning successfully. This account is necessarily speculative, since measuring innovation and redundancy at the item level will not capture a wide range of phenomena, such as swapping the orientation of axes or modification of the counting conventions, which occurs frequently in the maze game.

Perhaps the simplest potential explanation for the apparent contrast between the local and global effects on semantic co-ordination is that the CRs undermine participants’ confidence in the interaction as a whole. So, although they are locally consistent in their response to the CR, they subsequently become more generally conservative in their choice of description types. If editing of the target turn is treated as an index of confidence prior to the CR then there is some support for this in the data. Concrete responses are more likely after CRs to an ‘edited’ target. This is consistent with a view that the CR aggravates the lack of confidence. A ‘confidence’ explanation, however, still provides no mechanism that can explain the trend towards more abstract forms of semantic co-ordination.

The data suggest that conceiving of clarification subdialogues as consisting of a local (minimally) triadic exchange as in existing accounts obscures their role in the global development of co-ordination. Hypothetically, in providing positive evidence of understanding, the automatically generated acknowledgments should constitute the end of the CR sequence, allowing the conversation to proceed at the same level as prior to the intervention. However, both the global semantic effect and the different patterns of reuse work against this delineation of local context, and are also not captured by the IM’s global/local distinction provided by cumulative exposure

through input-output matching.

This raises two important methodological issues. First, the experiment was designed to produce CRs with a frequency similar to everyday conversation. In practice this resulted in each participant being exposed to one CR approximately every 40 turns, a figure which is far greater than the frequency of naturally occurring CRs in the baseline group from the first experiment. If the effects of the CRs are global rather than local it is possible that this causes a ‘cascading’ of errors in which each successive CR undermines the confidence of the response given to the previous one. More importantly, it is also possible that the different CR types interfered with each other. In particular, it suggests that combining the ‘What’ and ‘Frag’ manipulations in a single within-subjects condition is problematic.

The second issue concerns the use of the four description types as the unit of analysis to index levels of semantic co-ordination. Within each type there is considerable variation in how the descriptions are constructed and used. As a result, the local ‘sub-description type’ changes in response recorded by the innovation and redundancy measures are not detected. The global vs. local contrast in the data could thus be an artifact of the measures of semantic co-ordination. This will be dealt with in more detail in the following chapter.

Chapter 8

Discussion

Why does reducing the opportunities for direct interaction retard the development of semantic co-ordination? The first experiment provides evidence that interfering with sequential coherence increases the communicative effort required by participants to question and clarify each other's turns, impeding the development of more abstract descriptions. The second experiment, which also manipulated the level of feedback given to participants, in this case through assigning levels of participation, demonstrates a lack of convergence in dialogue involving less participation. In contrast, active participation leads to the use of more abstract descriptions. Taken together, these experiments provide convergent evidence that opportunities for direct interaction, in particular the opportunities that participants have for engaging in cycles of repair, affect their ability to converge on more abstract descriptions. While the first two experiments interfered with the use of both positive and negative feedback, the third experiment suggests that it is the role of negative feedback, in particular participants' ability to question and clarify each other's location descriptions that is fundamental to the development of more abstract descriptions.

In section 3.7 it was suggested that the co-ordinated use of abstract descriptions is indicative of higher semantic co-ordination: unlike concrete descriptions, which draw on immediate features of the particular maze being solved and are used in a more ad-hoc fashion, abstract descriptions draw on commonalities between mazes, allowing generalizations and more specific and more systematic comparisons between location descriptions. The data from the experiments confirm the role of interaction in the development of more abstract descriptions in becoming more semantically co-ordinated. Participants in the maze game do not attain this level of co-

ordination through explicit negotiation (Pickering and Garrod, 2004) nor through autonomous (Clark, 1996b) choice from a set of pre-existing descriptions, rather these descriptions emerge out of interaction (Healey, 2004). For example, although it might be assumed that ‘primitives’ such as ‘box’ or ‘square’ might be thought to be commonly understood in a population of English speakers, their use frequently poses problems, such as whether to count the links between boxes, whether to include non-instantiated boxes or whether to count from 0 or 1. Resolution of the ambiguities surrounding the use of ‘primitives’ in the maze game is a pre-requisite for convergence on the more abstract description types, which require even greater co-ordination between participants for their successful use.

The results from experiment three present a paradox concerning the role of local and global context in the resolution of ambiguities. Experiment one appeared to confirm existing accounts of clarification requests which conceive of participants’ ability to clarify and repair each other’s location descriptions as a local process that draws on the immediate context, allowing the quick resolution of communicative difficulties. However, description types only showed a global difference with a greater prevalence of concrete descriptions than in dialogue with no experimental manipulation. Locally, although participants took reliably longer to formulate responses to CRs with lower specificity, the responses were overwhelmingly of the same description type as the turn being clarified. As noted, this could be an artifact of using Figural, Path, Line and Matrix descriptions as the unit of analysis, since lower-level analysis of these descriptions showed significantly different patterns of responses, in terms of description constituents, to clarification requests with different strengths of specificity.

The focus of the interactive alignment model is on how participants come to use the same description types: the local mirrors the global, with turn-by-turn input-output matching resulting in global alignment. As described in section 3.7 this does not provide a full account of the shift toward more abstract descriptions observed in the original maze game experiments, due to the conservative nature of the model and its semantic neutrality. From this vantage-point, the global shift toward more abstract descriptions is achieved not through alignment but by an exploration process in which participants eventually realize that the Matrix descriptions are best suited to the task (Garrod, 1999).

8.1 Tacit modification against a backdrop of alignment

The proposal developed here is that alignment is actually the backdrop against which subtle, tacit changes are made in the process of developing more abstract description types. Convergence is not simply a cumulative effect of exposure, but is enabled by interactive repair operating at the constituent level allowing convergence on abstract descriptions through the exploitation of differences in sequential implicativeness afforded by the use of these elements. The four key elements of this claim are that: 1. The observed shift toward abstraction is not adequately explained at the level of description types; 2. Within types there is wide variation in the elements being used; 3. The use of these elements is more ad-hoc, local and not strongly tied to conceptualizations of the whole maze; 4. This variation is essential to the process of abstraction.

To illustrate this view, this section focuses on participants' use of counting, which involves the development of referring expressions that progressively refine the elements of the maze into higher and lower-order structures and their associated counting schemes. These examples are representative of a large class of naturally occurring CRs across the three experiments. They also serve to sidestep some problems with operationalizing 'abstraction'. In the maze game, counting is a fundamental step in the development of abstract descriptions, since they necessarily involve treating more than one feature of the maze as belonging to a category of the same type.

8.1.1 Granularity of response to CRs

Returning to the original maze game experiments (Garrod et al., 1984; Garrod and Anderson, 1987), the four description types are associated with distinct mental models that involve different features of the maze, as well as differences in the relationships between these features. An example of the former is whether the links between boxes are included in the model. An example of the latter is that in Line descriptions the individual squares which are the constituent elements of rows, can be composed into horizontal, vertical or diagonal elements. These claims are supported by the observation that the patterns of co-occurrence of particular items in participants' referring expressions correspond to the four kinds of description type and these four types generalize over a wide variety of possible location descriptions (Garrod and Anderson, 1987).

This raises the question of the granularity required in order to capture the role of the mechanisms involved in clarification dialogues and their effect on the development of semantic co-

ordination. This is exemplified by the two excerpts below¹.

- (8.1) | (1) A: I'm on the second row third on the left
 | (2) B: what?
 | (3) A: on the left
 | (4) B: ok move to the top left bit then
 | (5) A: cant get there
 | (6) B: where are you now?
 | (7) A: on the right
 | (8) B: where shall i move to now?
 | (9) A: from there go up two down three

In 8.1 it is not straightforward whether to classify A's response in (3) as a separate location description, or merely as a repeat of (1) with the deletion of some elements. This highlights the point that even triadic exchanges cannot be properly understood in isolation from their sequential context, since in the ensuing exchange, the participants resort to Figural and Path descriptions: A's subsequent use of *on the right* in (7) suggests that the response to the CR could also be conceived of as a concrete description. However, if it is conceived of as Figural, then this also obscures the relationship between the target turn and the response, rendering it equivalent to:

- (8.2) | (1) A: I'm on second row third on the left
 | (2) B: what?
 | (3) A: on the right sticking out bit

This highlights the problem of treating maze game turns as consisting of tokens of the four types, since as the excerpts above show, establishing the boundaries between individual descriptions is problematic (cf. Taylor and Cameron, 1987). This is of importance, since in the first instance, participants do not seem to be troubled by these boundaries. Instead, as suggested by the measures of redundancy and innovation in the third experiment, it appears that the co-ordination processes underlying the shift toward more abstract descriptions operate at the constituent level. Focusing on the changes made in response to the CRs highlights the observation that between turns (1) and (3), A, in redoing the location description does not reuse *row*, dropping an element of the description that involves counting. When A resumes counting, it is in using a Path description, which requires weaker co-ordination of counting schemes, as there is no need to establish axes or an origin.

¹The maze is simplified to facilitate illustration: switches, gates and goals have been removed as this maze depicts a superposition of both participants' mazes.




Figure 8.1: Maze corresponding to excerpt 8.1. Numbers in parentheses correspond to line numbers of excerpt.

Within existing models there is the expectation that after a clarification sequence, the conversation continues at the same ‘level’ (Clark, 1996b). If the analysis is performed with such a ‘ladder model’, this does appear to be the case. However, looking at the semantics of the descriptions used suggests a different story. For example, excerpt 8.1 shows that the effects persist after the CR, with participants continuing the use of concrete descriptions. From one perspective that brings the development of abstract descriptions into the foreground, the effect of the CR stretches until the hiatus involving the use of less abstract descriptions is followed by the resumption of the dialogue with more abstract descriptions. The argument here is that these modifications (additions, substitutions and deletions of constituents) which are made when participants encounter communicative difficulties are fundamental to the migration toward abstraction.

8.1.2 Variation within description types

The description type schema does not deal consistently with hybrid descriptions. For example, Path descriptions involve the use of another, usually Figural, starting point for describing the subsequent traversal through the maze. The other description types do not allow hybrids, although Line descriptions often bear more resemblance to concrete descriptions, for example: *I’m on the left of the row of four* relies on particular salient features of the maze, and is arguably

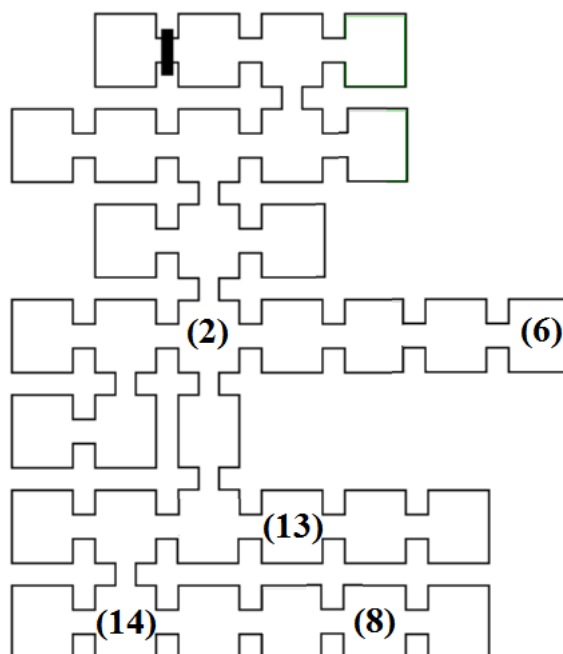


Figure 8.2: Maze corresponding to excerpt 8.3. Numbers in parentheses correspond to line numbers of excerpt.

less abstract than its canonical equivalent, *5th row 1st square*. Within Matrix description types, the counting scheme is frequently introduced through explicit negotiation and then becomes progressively abbreviated: *Columns are up, rows are down; column 5 row 1; col 5 row 1; c4 r 5; 4 5*. These are not always straightforward contractions of referring expressions as observed in standard referential tasks (Krauss and Weinheimer, 1966), as some of these stages occasionally involve changing the counting scheme to one previously used on another domain of the maze, or negotiate a counting scheme anew. For example, some pairs resorted to using a different origin when changing to a scheme involving a letter and a numeral, such as using $(7,1)$ to refer to the same location as *A1*.

Obviously the use of any classification schema is a trade-off, however there are generalizations over stretches of dialogue that are not reflected by changes within description types as described above. A particularly clear example of this is provided by the following excerpt. See Figure 8.2 for illustration of maze configuration ².

²The maze is simplified to facilitate illustration.

- (8.3) (1) B: wheres your goal
 (2) A: starting at top left three down two right
 (3) A: ok, so where r u?
 (4) B: in the top bit above my goal
 (5) A: wheres ur switch?
 (6) B: next to that right at the end of the arm thing
 (7) B: where u need to get to?
 (8) A: to bottom arm, in from right
 (9) B: Ummm
 (10) B: can only get to the second arm
 (11) A: second?
 (12) B: down
 (13) A: k, can only reach fourth from left arm second down
 (14) B: gates open am in 3rd arm, 2nd square

This excerpt shows participants developing the use of the term *arm* to refer to the protruding elements. It is first introduced by B in (6) as a concrete description as the *the arm thing*, and is reused by A in (8) to refer to another location with the additional qualifier *bottom arm*. Importantly, B's subsequent use of *arm* in (10) does not refer to the protruding element as *the middle arm*, as might be expected. Instead, B refers to it as the *second arm*, which appears to cause problems for A who initiates repair. The mutually ratified outcome of the clarification sequence suggests that the problem encountered by A might be due to ambiguity concerning whether *second* referred to the order in which the protruding elements were introduced to the conversation, in which case the second arm would be the lowest, or whether *second* was being used to count the protruding elements. Alternatively, A's CR can be seen to be clarifying whether to count from the top or the bottom. However, since there are three elements, the second element should in both cases be the one in the middle.

This dialogue shows evidence of participants co-ordinating on a more abstract description, but neither the shift observed in this excerpt nor the local clarification request that seeks to resolve the counting scheme on which part of this shift hinges are adequately captured at the level of description schema. The initial introduction of *arm* is strongly concrete, but over the ensuing turns, is systematically generalized. A counting scheme develops, starting from the top of the maze, suggesting that the counting scheme used in the second location description is reused to count over the protruding elements. Further, the participants also develop a counting scheme for the lower-order elements of the *arm*, adopting a scheme that counts from the rightmost node. This separation of maze elements into higher and lower-order structures combined with the development of counting schemes for both are the hallmarks of abstract descriptions. In fact, the location

description *3rd arm 2nd square* bears more resemblance to Line or Matrix descriptions, except that it does not involve a full model of the entire maze, and can only be applied to mazes that have salient horizontally protruding elements. Importantly, this excerpt shows the rapid development of abstraction, through the interplay of several different processes involved in the collaborative co-construction of more specific location descriptions, highlighting the simple but important process of the parallel development of counting schemes for higher-order and lower-order elements in the shift toward abstraction.

As excerpt 8.3 shows, the development of abstract descriptions does not necessarily mean that they extend over the whole maze. For example, participants frequently use row descriptions that are restricted to a subset of maze locations, such as the top half of the maze, referring to the other locations with other description types. Participants also frequently use Path descriptions when describing locations that are close to salient features of the maze such as *Two away from the goal* (Garrod and Doherty, 1994). At the item level, the elements used are not restricted to description types. Concrete elements such as *sticking out* or *left half* are both used as single objects of reference, as starting points for Path descriptions, as origins for abstract descriptions, and can also be further reified into countable elements, the counting schemes for which are frequently drawn from the counting of other elements. It is notable that the use of elements as *at the top* or *on the left or right hand side* are to be found in all location descriptions, and are especially prevalent in clarification subdialogues.

This wide variety of ambiguous elements that are constantly being recombined with each other suggests that convergence on a location description involves a far richer semantic space that is, at least initially, more fragmented, and more ad-hoc and ‘local’ than is captured by the schema of description types.

8.1.3 The role of variation in interactive repair

The role of variation in the use of description types as a factor in co-ordination is identified in Garrod and Doherty’s (1994) experiments on group co-ordination (see section 3.6.2). They use this to explain why a group that collectively encounters and resolves more co-ordination problems (and hence more variation in description types) develops more stable abstract descriptions.

The present results confirm the importance of opportunities for engaging in clarification requests for co-ordination. However, they pose a problem about the interaction between local and global context, since the global effect on descriptions is not reflected in participants’ local

changes of description type in response to clarification. It is not the case that on encountering difficulties participants immediately switch to a different description type. Instead, co-ordination processes operate on the constituent elements of the location descriptions.

The CM and the IM are not readily adaptable to giving an account of these modifications in response to clarification requests. If a misunderstanding occurs between levels, the CM places constraints on the kinds of repair performed in its resolution. Within levels, there is the general expectation that some kind of reformulation, specification or elaboration should take place, but no mechanisms are given that might constrain what form this should take. Similarly, Garrod and Doherty (1994) predict that on encountering trouble, participants should shift to the most frequently primed prior description type, yet the local changes in response to difficulty do not in the first instance occur at this level.

Across the three experiments, one frequently observed pattern in naturally occurring and artificially induced clarification sequences was the addition, deletion or modification of elements from the target turn and the preceding dialogue that appear to facilitate a reduction in the level of abstraction in the ensuing dialogue. The choice of modification does not appear to be arbitrary as the dialogue trajectories engendered by this modification predominantly involve a brief hiatus in which more concrete elements are used. After the initial introduction or reuse of these concrete elements, globally the CRs have the effect of allowing the co-ordinated continuation of the use of abstract descriptions. Across the three experiments, there are far more clarification sequences of the form

(8.4) | (1) A: Abstract description
 (2) B: Clarification request
 (3) A: Concrete response
 (4) B: Abstract description

than

(8.5) | (1) A: Concrete description
 (2) B: Clarification request
 (3) A: Abstract response
 (4) B: Concrete description

Although there are some clear cases where the modification of these elements corresponds to the boundaries between description types, this modification occurs at the constituent level. This is illustrated by excerpt 8.6 below ³.

³Goals, gates and switches have been removed to simplify the maze.

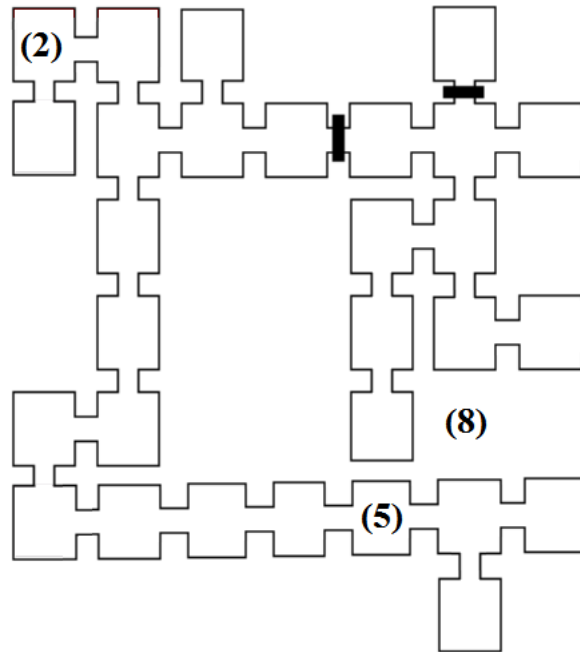


Figure 8.3: Maze corresponding to excerpt 8.6. Numbers in parentheses correspond to line numbers of excerpt.

- | | | | |
|-------|------|----|--|
| | (1) | A: | k, so am two spaces to the left |
| | (2) | B: | yeh ok am 1st row leftmost |
| | (3) | B: | ok |
| | (4) | A: | ok |
| | (5) | A: | go 6th row, fifth box |
| | (6) | B: | fifth?? |
| | (7) | A: | yeah |
| | (8) | B: | that dont exist, u sure we got the same mazes? |
| (8.6) | (9) | A: | should have countin 12345 |
| | (10) | B: | what? |
| | (11) | A: | 6th row fifth box from left hand side top |
| | (12) | B: | man that dont exist, only have two rows on right hand side |
| | (13) | A: | AHHHH, where you counting from with rows, you mean along or up down? |
| | (14) | A: | You never used a spreadsheet? |
| | (15) | B: | ... |
| | (16) | A: | OK, so am in 6th row 5th column, you got it? |
| | (17) | B: | Yeah |

This excerpt shows participants experiencing confusion surrounding the use of the term *row*. Initially, A uses *row* to denote horizontal vectors in the maze, and B as vertical vectors. In this excerpt, the participants shift from Line descriptions to a single Matrix description that combines both participants' axes used for the original Line descriptions. The first CR in (6), made by B,

fifth? is confirmed with a *yeah* that does not resolve the difficulty. The origin of the confusion is that B's original utterance *yeh ok am 1st row leftmost* is actually describing a location within the leftmost column, and is effectively explicating the counting scheme of counting columns from the left, hence B is in the leftmost column which is also the first column. This confusion persists after the first attempt at clarification and is exacerbated by A enumerating spaces in (9) that do not appear to B to exist on the maze. A's response (11) to B's subsequent *what?* inserts the new elements *from top left hand side top* into the original location description. This attempt is unsuccessful, since this description is also taken by B to mean the 6th vertical row. Importantly, in (12), B adapts A's introduction with *right hand side* and changes from ordinal to nominal counting of rows. This allows A to detect the confusion surrounding the use of axes, as this maze configuration has three horizontal rows that extend to the right edge of the maze window. Arguably, B's shift from ordinal to nominal counting of rows as well as A explicitly counting squares with *12345* are also evidence of relying more on the salient features of the maze, instead of an abstract schema of countable vectors. So, whilst in a sense the dialogue becomes more specific, narrowing down on the kind of counting problem, the elements used to accomplish this become progressively less abstract, shifting from an abstract description to counting individual maze locations, to using *left hand side* which finally helps in resolving the difficulty.

Excerpt 8.6 highlights how participants reuse elements from the preceding dialogue to resolve misunderstandings. Importantly, participants reuse *each other's* descriptions. Although the participants are aligned at the level of description types, the elements introduced and reused with each subsequent attempt at clarification allow a form of 'triangulation' on the nature of the problem. Importantly, this excerpt highlights the interaction between local and global effects of the CRs on semantic co-ordination. Locally, in response to clarification, more concrete elements are introduced, but after this brief hiatus the overall effect on the dialogue is an increase in abstraction.

This 'exploration' process (Garrod, 1999) is not adequately modelled as a 'landscape' of co-ordination equilibria. Rather, it involves systematic use of more concrete constituents at the sub-description level to resolve differences in interpretation.

The word-level measures of innovation and redundancy show that the participants draw on a wide variety of elements from the preceding dialogue in attempting to resolve communicative difficulties. The rich variety in the patterns of reuse and their effect on the subsequent dialogue

suggest that a turn-by-turn analysis of alignment provided by entrainment scores is insufficient, since what these modifications highlight is the importance of *differences* between turns and their effect on the ensuing dialogue. This is illustrated by the excerpt below (see Figure 8.4 for maze corresponding to transcript)⁴.

- (8.7)
- (1) A: My switchpoint is in the 1st line where are your switches?
 - (2) B: hey, switchpoint is 2 in third row
 - (3) A: 2?
 - (4) B: from left bit
 - (5) A: have your gates opened?
 - (6) B: No
 - (7) A: am on switchpoint three down and four left from the side bit
 - (8) B: No ur not, shld be open.
 - (9) A: Am 2nd square of the third row
 - (10) B: Well it's not open
 - (11) B: You sure you're in the 2nd square from left side ?
 - (12) A: Give me directions from the top right of the two columns
 - (13) B: go three down 2 left.
 - (14) A: what?
 - (15) B: from top bit go three down
 - (16) A: one above bit on its own?
 - (17) B: Yeah then then 2 left
 - (18) A: 2?
 - (19) B: squares left yeah, hurry am getting hungry
 - (20) A: open?
 - (21) B: Yeah. Where's do u want me to go
 - (22) A: 6th column 4 down
 - (23) B: ???????
 - (24) A: ??
 - (25) A: 6th column frm right bit two down
 - (26) B: open?

This excerpt starts where A introduces the first Line description into the dialogue. Immediately, B also introduces a Line description into the dialogue, although using *row* instead of *line*, which is queried by A. It appears that A's difficulty in understanding B's Line description stems from the use of *2 in*: this particular row is only five squares wide, with a gap of two squares separating it from the edge of the window containing the maze. On the basis of this description alone, it is ambiguous whether to count from the first square in the row or whether the *2 in* is actually describing the two non-instantiated squares constituting the gap. There is also the additional ambiguity of whether to count the links between the nodes in the maze or the nodes themselves.

⁴The maze is simplified to facilitate illustration.




Figure 8.4: Maze corresponding to excerpt 8.7. Numbers in parentheses correspond to line numbers of excerpt.

There are therefore at least three possible locations within the third row that *2 in* could be referring to. B's response involves a deletion of *2 in third row*, adding instead the more concrete *from left bit*. Although the introduction of *bit* might be expected to favour an interpretation that the nodes should be counted from the leftmost edge of the row, as *bit* suggests it is referring to some object, as opposed to a gap between the edge of the window and the leftmost node of the row. In (7), A reuses the term *bit* by reformulating the current location as a Path description, although this does not resolve the ambiguity. Nor does adding *square* to assist enumerating the nodes and substituting *line* (1) with *row*. B's next attempt in (11) is effectively a repeat of the the response to the clarification, substituting *bit* for *side*. At this juncture, both participants are unclear about each other's (and possibly their own) use of both terms, which results in another attempt (12)-(20) which effectively serves as a triangulation on the problem surrounding how to count in the maze. (12) is also the first introduction of vertical vectors into the dialogue, although here it is introduced figuratively as a single entity comprising two columns. This Path description starting from the top of the two columns introduces a new problem, namely that of whether it is possible to count nodes that are not connected by a path: the second part of the Path description *two left* involves such a jump between two nodes; B's decomposition into two parts, the first part substitutes *right of the two columns* with the less specific *bit*. The second part also reuses *bit*, here

clearly referring to an actual node. This should resolve the original difficulty of (4), although at (18), A clarifies that it is possible to count across squares with no connecting nodes, which results in the introduction in (22) of a column description which arguably reuses the vertical counting convention used earlier (1) (2) to count the rows from the top, but this time to count the individual squares within the column, and the columns, which were originally introduced figuratively are now counted from the right hand side, in contrast to the squares within rows, which were counted from the left. A's response to B's open-class repair that signals problems understanding this novel introduction of a column description parallels the first CR in (2)-(4) yet here results in successful continuation of the dialogue, even though the original problem of determining how exactly to count *in* was not directly resolved and what counts as a *bit* was only tacitly co-ordinated on.

This excerpt shows another instance of participants' initial conflicting use of abstract descriptions being ratified through cycles of repair. Here too, at the item level, the choices made in responding to repair appear to follow a similar pattern: one of the ways in which abstract descriptions become established is that they are initially attempted and subsequently clarified with a brief hiatus where more concrete elements are introduced or reused from the previous dialogue, or the more abstract elements are substituted or deleted in response as in excerpts 8.1 and 8.6. This arguably facilitates participants' continuation of the interaction, allowing the finer details of the abstract descriptions to be settled tacitly before either continuing with the 'same' but fleshed out abstract description, or as in the excerpt above, with a different one which relies on the generalizations made possible by the hiatus involving these concrete elements.

The two excerpts show how the semantic effects of the CRs stretch beyond the immediate subdialogue and highlight the complexity in the patterns of reuse of elements in CR sequences. The addition of elements in CR responses is treated in the CM as the specification of more information, but this ignores what is occurring at the constituent level. Existing CR models of Purver et al. (2003) or Schlangen (2004) use notions of reformulation or elaboration that do not capture these phenomena. Arguably it is these modifications at the constituent level and the subsequent modification by the initiator of the CR response (as in excerpt 8.1) that determine whether the utterance is an elaboration or specification.

In addition, examples such as 8.1 which only delete elements from the target turn occur regularly across the three experiments. These present a problem for existing models as the responses provide no additional information, yet still have the effect of resolving the misunderstanding.

This contrasts with Clark and Marshall's (1981, p. 48) expectation that to be effective, repair 'must add or alter descriptors, but not delete them'. Importantly, these CR responses that simply delete elements from the target turn are not treated as repeats and queried again, but appear to invoke different trajectories than if the target turn had been repeated verbatim.

Although there does appear to be a large class of these kinds of reparative sequences, this is not the only way in which CRs are deployed in the maze game. However, there is a cumulative body of evidence that this deployment is not done explicitly, but is done tacitly. The strength of the IM over the collaborative model is that it recognizes this. However, the priming mechanism of input-output matching appears to be more the product rather than the precondition of successful semantic co-ordination.

8.1.4 Summary

Bringing the role of repair in semantic co-ordination into the foreground provides an account that allows for participants to start with radically different models of the maze. In the interactive alignment model it is assumed that alignment occurs due to priming operating on models that both parties to the conversation possess prior to the conversation. It does not provide a satisfactory account of how these models emerged in the first place, nor how very different models could be able to converge through priming. In contrast to the prioritization of input-output matching which is used to account for the stability of pre-existing models of the maze, the results of the three experiments strongly indicate that interactive repair plays a key role in this process. Semantic co-ordination cannot be simply explained as a cumulative effect of exposure: in the original maze game experiments, participants in the sub-community group initially used more Line descriptions before converging on Matrix descriptions. The same pattern was observed in the second experiment, and in the first experiment, although participants were exposed to the same number of turns, in dialogue where the use of repair was inhibited, the participants used more concrete descriptions.

Instead of focusing on what is kept 'the same' across turns, it seems that variation in description types is a key communicative resource for performing repair in narrowing down representations of the maze. Whilst across turns there is high level matching of description types, this is exploited for the modification of low level elements: one kind of trajectory outlined above involves 'redoing' the same description that is ostensibly of the 'same' type, yet the simple addition of more concrete elements or the deletion of more 'abstract' elements does not merely provide

additional information to assist disambiguation. These modifications also bring about the possibility of ensuing dialogue trajectories that are less ‘abstract’. Importantly, the results of the first experiment strongly suggest that it is participants’ ability to modify *each other’s* representations that is key to the development of semantic co-ordination.

This raises a methodological point. Measures of entrainment, whether at the level of description type or its constituent elements are insufficient for measuring the development of semantic co-ordination. The CR excerpts given above demonstrate the importance of the relationship between what is kept ‘the same’ and what is modified. This also implies that the measures of innovation and redundancy employed in the third experiment are also inadequate for the task. What is required is a measure that is sensitive to the kinds of description being used and the kind of repair being performed.

The account offered here suggests a different view of how abstract descriptions come to be established in maze game dialogue. It does not appear from the data that as a general rule abstract descriptions are introduced into the dialogue only when participants have solved enough co-ordination problems that would allow them to adopt abstract descriptions in a trouble-free way. Unlike Wittgenstein’s example of a pupil moving from one set of mathematical problems to the next once an expert has ascertained the pupil’s skill, in the maze game, participants frequently introduce abstract descriptions without having fully worked out the counting schemes for both higher and lower-order elements (cf. Larsson, 2005). Thus it is not the case that participants follow a plan of methodically setting out what is required in order for them to get the abstract description schemes off the ground. Instead, the descriptions emerge opportunistically (Clark, 1996b) and the details of the referring expressions are worked out through cycles of repair. This frequently results in an abstract description that is different from the original due to the interactional trajectories engendered by the addition, substitution or deletion of its constituent elements. Convergence on abstract descriptions rests not on parity of representations but on asymmetries between participants that allow the mixing and reuse of the elements introduced by each other in order to ‘triangulate’ on the problems encountered in their use.

This works against using the occurrence of abstract descriptions as a general index of semantic co-ordination, since in dialogue with fewer opportunities for clarification, participants initially use more abstract descriptions. What is needed is a model that is sensitive to both the semantics of the descriptions in use and the kinds of clarification being performed. In contrast to exist-

ing accounts, where longer contribution cycles are seen as indexing difficulty in communicating, the first two experiments suggest the opposite. The clarification of initial use of abstract descriptions is in some senses indicative of higher co-ordination, since without them, misunderstandings concerning the use of abstract descriptions persist in the dialogue.

The excerpts above focused on two similar mechanisms of adding and substituting elements in reparative sequences. The complexity of the patterns of reuse show that conceiving of repair as a triadic sequence operating on description types is inadequate for an account of the development of semantic co-ordination. What is required is a more detailed analysis of the sequential trajectory in which repair is performed. To be sure, the mechanisms outlined above are not the only way in which participants use repair in the maze game, but they clearly illustrate how their use drives semantic co-ordination involving complex patterns of modification that are not adequately captured by existing models due to their prioritization of what remains the ‘same’.

The account developed here also emphasizes the importance of input-output matching, in this example, as in the excerpts above, the CR responses and the ensuing turns are reliant on the gross alignment it provides in order to facilitate the additions, substitutions and deletions that allow continuation of the dialogue. This, however, raises the issue about the relationship between lexical and semantic priming, as these patterns of reuse of elements in different descriptions against a backdrop of alignment suggest that this relationship cannot be simple reinforcement or percolation.

What would a non-conservative model of these phenomena look like? Such a model would necessarily have to take the semantics of the descriptions used into account, bringing the role of repair into the foreground. The data provided by the experiments all point toward the complex interplay between repair and the variety of descriptions being used in developing semantic co-ordination, requiring further experimental investigation of these phenomena.

8.2 Conclusions and further work

The account developed here provides guidelines for further investigation of these phenomena and has implications for existing models of communication. It emphasizes that it is not the matching of referring expressions or representations but the differences between elements which are of key importance to the development of co-ordination. This involves a critical interaction between low-level and high-level processes that occurs at the constituent level: similarity and

difference of description type is of less consequence than the lower-level interplay of constituents through addition, substitution and deletion. Further, these modifications involve the systematic exploitation of concrete, not abstract referring expressions.

As the excerpts above show, the matching that occurs between turns neglects what is changed, thus entrainment measures, whether at the item or description type level are inadequate for the task of investigating dialogue, and are also inadequate as an explanatory device to account for how participants come to develop semantic co-ordination. If, as is suggested, the variety of elements is of key importance to the development of semantic co-ordination, the question emerges of how this variety is narrowed down into successful use of abstract descriptions. The excerpts above showed clarification trajectories involving complex patterns of reuse that assisted participants in co-ordinating on location descriptions.

It appears from the experimental data that both introducing artificial clarification requests and blocking the opportunities for repair interfere with participants' ability to exploit the variety of items introduced by each other during the dialogue to resolve the difficulties encountered in co-ordination of abstract descriptions. Even though there are very few CR questions in the baseline condition of the first experiment, and even fewer CR questions in the dual window condition, the third experiment highlights their importance in the development of semantic co-ordination. It shows that even relatively few CRs can have a pronounced effect on the dialogue. The excerpts above showed the use of multiple CRs in succession that allowed the set of potential problems to be narrowed down. As noted by Schegloff (1992), interlocutors frequently display an acute sense of the kind of trouble that is signalled by initiation of repair. This highlights an additional problem concerning the experimental manipulation of introducing artificial clarification requests into dialogue: each artificially introduced CR not only indicates immediate local difficulty, but can also be taken to be signalling that a previous CR had not been adequately resolved, resulting in a 'cascading' of externalized inferencing (Pickering and Garrod, 2004). Possibly the secondary effect of this cascading is that it interferes with naturally occurring CRs, rendering them less effective. By hypothesis, as in the first experiment, this also interferes with participants *mutually modifying* (Healey et al., 2007) each other's turns in drawing on the rich variety of fragmentary and ad-hoc elements in order to deal with the communicative difficulties that arise when attempting to co-ordinate on more abstract schemes.

It is unclear to what extent the asymmetries caused by the chat server not relaying CR re-

sponses to the other participant contribute toward this effect. There are many possible contributing factors, such as disruption of turn-taking, or causing gross asymmetries in the common ground, for example by blocking a request for information or blocking the first mention of a location description. However, artificial CRs also introduce finer asymmetries, as participants do not always respond with a full location description. Instead, as exemplified by the excerpts above, participants add, substitute and delete elements from the target turn, and it is argued that these elements are generally more tied to more concrete features of the maze than to an abstract description. For example, CR responses frequently only repeat the elements that do not involve counting, which are then reused by the other participant in continuing the dialogue using more concrete descriptions. This kind of subsequent dialogue trajectory which is evident in naturally occurring clarification sequences in the maze game is prohibited by the chat tool for artificial CRs, as the modification of the target turn is not relayed to the other participant.

More importantly, the modifications that occur at the item level without changing description type suggest that the actual outcome of the CR on semantic co-ordination is strongly dependent on the ‘uptake’ by the CR initiator of the response. The reuse of more concrete elements implicates two possible trajectories: continuation with the description type of the target turn and a continuation, as in excerpt 8.1 or 8.6 line 12 with concrete descriptions. This then suggests another possible reason why there was a lack of a shift in descriptions in the third experiment.

A related issue is that in the third experiment, all responses to CRs were acknowledged with an ‘ok’, preventing the artificial clarification trajectory from unfolding any further. In the initial pilot studies for the CR experiment, the random number controlling the interval between interventions had no lower bound, in some instances participants’ responses to CRs were queried with a further CR. Although there were only 7 of these dual CRs, 4 of them resulted in participants using location descriptions of a different type according to Garrod et al’s schema. If, as the data strongly suggests, participants’ reuse of elements when faced with communicative difficulties is motivated by ‘concreteness’, performing a second CR on the response to the original CR that is sensitive to the modifications made between target turn and the response would allow a much more detailed understanding of regularities in the patterns of reuse of these elements in clarification requests. A more promising approach would be to selectively block participants’ use of different kinds of repair. This would not, as in the first experiment, directly interfere with other interactive exchange mechanisms and would allow better understanding of how repair contributes

to narrowing down variety in the development of semantic co-ordination.

One particular mechanism employed by participants to narrow down the descriptions used in the maze game is embedded repair: in excerpt 8.7 participants converged on the use of ‘row’ through the substitution of ‘line’. As described in section 3.3.1 embedded repair is only identifiable as such by the modification of a single element against a backdrop of alignment, allowing the substitution of the problematic element.

Instead of performing coarse-grained manipulations as in the first experiment that interfered with a broad range of interactive mechanisms, the chat tool affords the selective substitution of synonyms such as ‘square’ and ‘box’ or ‘row’ and ‘line’ which permits the blocking of participants’ use of embedded repair through both enforced alignment and misalignment. This would allow closer understanding of the relationship between alignment and the exploitation of differences in levels of abstraction by the experimental separation of both phenomena. The experimental data all point towards this exploitation of differences occurring at the sub-constituent level, there being a critical interaction between high-level matching and low-level addition, substitution and deletion of constituents. The question that emerges is what governs the choice of addition, deletion or substitution? Why and how much should be modified?

The chat tool offers the opportunity of investigating these phenomena through context-sensitive manipulation of participants’ opportunities for repair-initiation. This would allow teasing apart the possible influences of levels of grounding, effects of recency of use, whether an expression had been introduced explicitly, whether it had been used successfully or not, and how these influences interact with the kind of item being clarified. This would go beyond Purver et al’s (2003) investigation of the differences between content and function words and effect of first vs. subsequent mention. An understanding of these effects would lay the groundwork for investigating the relationship between high-level matching and low-level modification: the constrained vocabulary of maze game locations permits the manipulation of enforced alignment at the higher level in conjunction with enforced misalignment at the lower level through systematic translation of location descriptions, thereby creating controlled asymmetries that would lead towards establishing generalizations in the patterns of modification that occur in clarification sequences.

The immediate question that emerges is the extent to which these findings generalize beyond the maze game to different modalities or dialogue in general. The maze game setup enforces a communicative context which presents participants with problems of co-ordination on a

structured relational domain that requires the integration of individual perspectives. These three features are what Schwartz (1995) sees as promoting the emergence of abstract descriptions as a direct consequence of encountering and resolving co-ordination problems. One experiment conducted by Schwartz involved participants attempting to solve the direction of rotation (clockwise/anticlockwise) of a sequence of interconnected gears. Schwartz observed that in developing an abstract 'parity' rule participants often co-ordinated tacitly on more concrete gestural simulations of the rotating gears before developing a counting procedure which allowed them to establish that odd-numbered sequences of gears rotated anti-clockwise. Another task required participants to develop a representation of interdependencies between animals and their habitats. Here too, dyads incorporated elements of both members' representations into more abstract formalisms. In these experiments, Schwartz consistently found that the representations that emerged collaboratively through direct interaction between participants differed qualitatively from those used by solitary participants attempting the same task. This difference is especially clear in a set of further experiments (Schwartz, 1999; Schwartz and Black, 1999) where participants were presented with containers of equal height but different diameters that had a line drawn at the same height from the base to represent an imaginary water-level. Participants were set the task of imagining the angle of tilt at which the water would pour from both containers in order to determine whether both containers would pour at the same angle or whether one would pour before the other. In this task no dyads arrived at the correct answer, compared with 100% of individuals who were instructed to close their eyes and imagine tilting both containers. In this case, the collaborative process led to unproductive abstract descriptions. This reinforces the argument made earlier that the mere presence of abstract descriptions does not necessarily index full co-ordination, and underscores the importance assigned by this thesis to collaborative repair mechanisms that address and resolve the kinds of problems encountered in the development of abstract descriptions.

In addition, experiments conducted by Healey, Swoboda et al.(2007; 2002) suggest that the findings of this thesis also extend to the modality of graphical interaction. Participants were assigned the task of communicating via a shared interactive whiteboard in order to establish whether they were listening to the same musical piece. In this task, participants' drawings could be reliably classified into concrete and abstract descriptions. The former drew on impressions such as landscapes or facial expressions, and the latter drew on features such as pitch, volume or

tempo. In order to reduce the opportunities of participants signalling difficulties with elements of each other's drawings through circling, underlining or drawing arrows, participants' whiteboards were divided into separate windows similar to the dual window chat tool of the first experiment. This proved to have a reliable effect of reducing participants' use of abstract drawings, paralleling the findings of this thesis that it is participants' ability to modify and locate problems in each other's representations that is key to the development of abstraction. Further, assigning participants to different sub-communities led to the development of community-specific representations. When members of different sub-communities subsequently interacted with each other, they relied on more concrete descriptions in resolving communicative difficulties. Here too, the results demonstrated the systematic exploitation of concrete rather than abstract elements on encountering difficulties.

This body of evidence provided by experiments across modalities strongly suggest that the phenomena described in this thesis are of importance to accounts of language that seek to explain the development of co-ordination. This leads to the question of the relationship between the patterns of modification in problematic and unproblematic dialogue as participants also incorporate each other's descriptions in dialogue where no overt problems are signalled. In the interactive alignment model this is seen as the main locus of co-ordination, as it is here that priming is most effective. In the collaborative model too, as part of successful referring, participants both incorporate each other's descriptions and also delete elements on successive use. The account developed here reverses these priorities, bringing the role of miscommunication into the foreground, since from this perspective it is when the convention doesn't 'work' as expected that participants gain a sense of its applicability. The chat tool could be used to establish whether the tacit exploitation of more concrete elements in problematic dialogue functions in the same manner if introduced into dialogue with no explicit signalling of communicative difficulties. This would allow a much deeper understanding of the extent of participants' orientation toward levels of abstraction, since it could be established whether shifts toward more concrete descriptions in the absence of any overt signalling of repair are taken as signalling communicative difficulties. This would allow refinement of Brennan and Clark's (1996) notion of conceptual pacts, since the results show that participants do not treat description types equally.

The other question that arises is where the boundary between the provision of additional information and the different sequential trajectories engendered by these tacit modifications lies.

This could be addressed by further experiments involving the selective capturing of repair initiation of different levels of specificity and translating them into more explicit requests for information. This would prevent some of the problems caused by the introduction of artificial clarification requests, providing a direct comparison between the two trajectories of blocked vs. rendered explicit and allow an understanding of the parameters within which the clarification sequences operate.

As described above, the response of the CR initiator to the modification of the target turn is key to determining whether the dialogue continues with the use of concrete descriptions. The chat tool offers the capability of identifying the clarification requests and modifying the response, engaging the participants in artificial clarification subdialogues in which the server responds to the CR instead of generating it. This would complement the approach of rendering the clarification requests more explicit, as it would allow the context-sensitive manipulation of responses, in order to investigate the second part of the triadic clarification sequences.

To conclude, the phenomena investigated in this thesis concerning the systematic exploitation of concreteness in responding to communicative difficulties are not fully addressed by existing models of dialogue. Formal models, being concerned with the update of information-states are not sensitive to these semantic differences. C.A.'s primary concern with demonstrating participants' orientation towards the mechanisms that are C.A.'s objects of analysis and its focus on the surface-structure of dialogue glosses over the semantic differences between concrete and abstract descriptions. The conversation analytic approach adds a methodological hurdle to providing an account of the tacit modifications made in dialogue, as one of their main hallmarks is that they are achieved tacitly, hence the difficulty in ascertaining interlocutors' orientation towards them.

The account developed in this thesis goes beyond C.A. in demonstrating that there is an interface between the procedural characterisation of repair and local and global semantic change. The experimental approach adopted goes beyond corpus analysis to demonstrate causal relationships between the procedural mechanisms of C.A. and the semantics in use. While the experimental investigation using the chat tool loses some sensitivity when compared with C.A.'s approach that typically involves the close analysis of a small set of examples, this loss of sensitivity is compensated by greater generalization. It also directly counters C.A.'s reservations concerning the experimental investigation of repair, as this thesis demonstrates that the use of the chat tool to manipulate the development of constrained language games is a fruitful and promising approach

to the investigation of the development of co-ordination.

These phenomena are of prime importance to theories of language use and evolution, in particular evolutionary simulations used to model the dynamics of language, as the development of abstract conventions presents an even greater bootstrapping problem. While Ginzburg and Macura (2005; 2006) address the gap in existing evolutionary models by incorporating clarification sequences into their agents' communicative protocols, which has a stabilizing effect on the language of the population, the semantics of the agents' language, as in Steel's simulations (2003; 2005) do not address the issue of the development of abstraction. This suggests a possible avenue for exploration and would also allow the potential to provide insights into the impact of these mechanisms on group dynamics.

The phenomena described in the first experiment are of immediate practical concern to the design of chat tool interfaces. The dual window setup is equivalent to existing implementations (UNIX chat). Although this has generally been superseded by instant messenger programs which use a single window, many online games, played by millions of subscribers, use a 'speech bubble' chat interface which is analogous to the dual window approach of separating speakers' turns. One of the key features of these games is that they require players to interact with each other about their virtual environment, thereby developing idiosyncratic conventions. The data presented here show clear effects of the interface on the kinds of expression used, strongly indicating that choice of chat tool should not rest purely on aesthetic grounds. Further experimentation is required to establish the effects of, *inter alia*, window size, number of interlocutors and 'scrolling' on co-ordination.

These phenomena are also of immediate practical concern to the design of dialogue management systems. The process of abstraction is rapid, often occurring in under 15 turns, suggesting that these phenomena might be operative in even relatively brief information-seeking dialogues. This has implications for handling both user-generated and system-generated CRs. Detecting these modifications in response to system-generated CRs would allow a much finer-grained indexing of co-ordination and more coherent selection of the next dialogue move or course of action. More importantly, the findings of this thesis suggest that mutual-intelligibility can also be sustained through the use of more concrete descriptions, allowing the exploitation of these phenomena by the system in responding to user-generated CRs, as the interactive phenomena (in particular the deletes) are not simply reducible to the provision of additional information.

The results obtained from the chat tool confirm the experimental approach adopted in this thesis. The data show that the phenomena are more fine-grained than the manipulations performed in the three experiments. However, the results strongly call for a further program of investigation. The chat tool provides an ideal platform for this investigation, as it offers the possibility of systematic, fine-grained and context-sensitive manipulation of unfolding dialogue, providing rich opportunities for further empirical work.

Bibliography

- Allwood, J. (1992). On dialogue cohesion. In *Gothenburg Papers in Theoretical Linguistics*. Dept of Linguistics Göteborg University.
- Allwood, J. (1995). An activity based approach to pragmatics. In *Abduction, Belief and Context in Dialogue, Studies in Computational Pragmatics*, Amsterdam. John Benjamins.
- Anderson, A. and Garrod, S. (1987). The dynamics of referential meaning in spontaneous dialogue. In *Communication Failure in Dialogue and Discourse*. Elsevier Science Publishers B.V.
- Anderson, J. F., Beard, F. K., and Walther, J. B. (2000). The local management of computer-mediated conversation. In *Herring (1999) Interactional coherence in CMC*.
- Atkinson, J. M. and Heritage, J. C., editors (1984). *Structures of social action: Studies in conversation analysis*. Cambridge University Press, Cambridge, England.
- Atkinson, M. and Drew, P. (1979). *Order In Court*. Macmillan, London.
- Austin, J. L. (1953). How to talk: Some simple ways. In *Proceedings of the Aristotelian Society*, volume 53, pages 227–246.
- Austin, J. L. (1962). *How to Do Things with Words*. Oxford University Press, Oxford.
- Austin, J. L. (1970). *Philosophical Papers*. Oxford University Press, London, 2nd edition.
- Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence on social perception and cognition. In Uleman, J. S. and Bargh, J. A., editors, *Unintended thoughts*, pages 3–51. Guilford Press.
- Barr, J. B. and Keysar, B. (2004). Is language processing different in dialogue? *Behavioural and Brain Sciences*, 27(2):22–23. Commentary on Pickering and Garrod (2004): Toward a mechanistic psychology of dialogue.

- Bavelas, J., Chovil, N., Lawrie, D., and Wade, L. (1992). Interactive gestures. *Discourse Processes*, 15:469–489.
- Boaz, K., Barr, D., Balin, J., and Paek, T. (1998). Definite reference and mutual knowledge: Process models of common ground in comprehension. *Journal of Memory and Language*, 39:1–20.
- Branigan, H. P., Pickering, M., and Cleland, A. A. (2000). Syntactic co-ordination in dialogue. *Cognition*, 75:B12–B25.
- Brennan, S. and Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 22:1482–1493.
- Cacciari, C. and Glucksberg, S. (1994). Understanding figurative language. In Gernsbacher, M. A., editor, *Handbook of Psycholinguistics*, pages 447–477. Academic Press, San Diego.
- Carletta, J., Garrod, S., and Fraser-Krauss, H. (1998). Communication and placement of authority in workplace groups - the consequences for innovation. *Small Group Research*, 29(5):531–559.
- Cherry, C. (1966). *On Human Communication: a review, a survey, and a criticism*. MIT Press.
- Clark, H. H. (1979). Responding to indirect speech acts. *Cognitive Psychology*, 11:430–477.
- Clark, H. H. (1985). Language use and language users. In Lindzey, G. and Aronson, E., editors, *Handbook of Social Psychology*, volume 2. Random House, New York, 3rd edition.
- Clark, H. H. (1996a). Communities, commonalities and communication. In Gumperz, J. J. and Levinson, S. C., editors, *Rethinking Linguistic Relativity*, pages 324–355. Cambridge University Press, Cambridge.
- Clark, H. H. (1996b). *Using Language*. Cambridge University Press, Cambridge.
- Clark, H. H. and Brennan, S. (1991). Grounding in communication. In *Perspectives on socially shared cognition*. APA Books, Washington.
- Clark, H. H. and Lucy, P. (1975). Inferring what was meant from what was said. *Journal of Verbal Learning and Verbal Behavior*, 14:56–72.

- Clark, H. H. and Marshall, C. (1981). Definite reference and mutual knowledge. In Webber, B. L., Joshi, A. K., and Sag, I. A., editors, *Elements of Discourse Processing*, pages 10–63. Cambridge University Press, New York.
- Clark, H. H. and Schaefer, E. (1989). Contributing to discourse. *Cognitive Science*, 13:259–294.
- Clark, H. H. and Schaefer, E. F. (1992). Dealing with overhearers. In Clark, H. H., editor, *Arenas of language use*, pages 176–197. University of Chicago, Chicago.
- Clark, H. H. and Schaefer, F. (1987). Concealing one’s meaning from overhearers. *Journal of Memory and Language*, 26:209–225.
- Clark, H. H., Schreuder, R., and Buttrick, S. (1983). Common ground and the understanding of demonstrative reference. *Journal of Verbal Learning and Verbal Behaviour*, 22:245–258.
- Clark, H. H. and Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 2:1–39.
- Cobley, P., editor (1996). *The Communication Theory Reader*. Routledge.
- Condon, S. (2001). Discourse *ok* revisited: Default organization in verbal interaction. *Journal of Pragmatics*, 33:491–513.
- Craig, R. (1999). Communication theory as a field. *Communication Theory*, 9:119–161.
- Drew, P. (1997). ‘open’ class repair initiators in response to sequential sources of troubles. *Journal of Pragmatics*, 28:69–101.
- Duranti, A. (1997). *Linguistic anthropology*. Cambridge University Press.
- Foucault, M. (1998). *Essential works of Foucault : trans Robert, J. Hurley*. New Press.
- Gabsdil, M. (2003). Clarification in spoken dialogue systems. In *Proceedings of the 2003 AAAI Spring Symposium Workshop on Natural Language Generation in Spoken and Written Dialogue*, Stanford, USA.
- Garcia, A. and Jacobs, J. B. (1999). The eyes of the beholder: Understanding the turn-taking system in quasi-synchronous computer-mediated communication. *Research on Language and Social Interaction*, 32(4):337–367.

- Garcia, A. C. and Jacobs, J. (1998). The interactional organization of cmc in the classroom. *Qualitative Sociology*, 21(3):299–317.
- Garfinkel, H. (1963). A conception of, and experiments with, 'trust' as a condition of stable concerted actions. In Harvey, O. J., editor, *Motivation and Social Interaction*. Ronald Press.
- Garfinkel, H. (1967). *Studies in Ethnomethodology*. Prentice-Hall, New Jersey.
- Garrod, D., Anderson, A., and Sanford, A. (1984). Semantic negotiation and the dynamics of conversational meaning. Technical Report 1, Glasgow Psychology Technical Report.
- Garrod, S. (1999). The challenge of dialogue for theories of language processing. In Garrod, S. and Pickering, M., editors, *Language Processing*. Hove: Psychology Press.
- Garrod, S. and Anderson, A. (1987). Saying what you mean in dialogue: a study in conceptual and semantic co-ordination. *Cognition*, 27:181–218.
- Garrod, S. and Doherty, G. (1994). Conversation, co-ordination and convention: an empirical investigation of how groups establish linguistic conventions. *Cognition*, 53:181–215.
- Gibbs, R. W. (1984). A critical examination of the contribution of literal meaning to understanding non-literal discourse. *Text*, 2:9–27.
- Gibbs, R. W. (2002). A new look at literal meaning in understanding what is said and what is implicated. *Journal of Pragmatics*, 34:457–486.
- Ginzburg, J. and Macura, Z. (2005). Lexical acquisition with and without meta-communication. In *The emergence and evolution of linguistic communication*. Springer.
- Goodwin, C. and Heritage, J. (1990). Conversation analysis. *Annual Review of Anthropology*, 19:283–307.
- Grice, H. P. (1957). Meaning. *Philosophical Review*, 66:377–388.
- Grice, H. P. (1969). Utterer's meaning and intentions. *Philosophical Review*, 78:147–177.
- Grice, H. P. (1975). Logic and conversation. In Cole, P. and Morgan, J. L., editors, *Syntax and semantics: Speech acts*. Academic, New York.
- Haapmaki, S. (2002). *Studier i svensk grammatikhistoria (Studies in the History of Swedish Grammars)*. Abon AkademiUniversity Press, Turku/Abo.

- Healey, P. G. T. (1997). Expertise or expert-ese? the emergence of task-oriented sub-languages. In *Proceedings of the 19th Annual Conference of the Cognitive Science Society*, Stanford, California.
- Healey, P. G. T. (2004). Dialogue in the degenerate case? Peer commentary on Pickering & Garrod (2004). *Behavioural and Brain Sciences*, 27(2):201.
- Healey, P. G. T. and Mills, G. (2006). Participation, precedence and co-ordination. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, Vancouver, Canada.
- Healey, P. G. T., Swoboda, N., Umata, I., and Katagiri, Y. (2002). Graphical representation in graphical dialogue. *Interactional Journal of Human Computer Studies*, 57(4):375–395. (Special issue on Interactive Graphical Communication).
- Healey, P. G. T., Swoboda, N., Umata, I., and King, J. (2007). Graphical language games: Interactional constraints on representational form. *Cognitive Science*, 31(2):285–309.
- Heritage, J. C. (1984a). Analysing news interviews: aspects of the production of talk for an overhearing audience. In *Handbook of discourse analysis*. Academic Press.
- Heritage, J. C. (1984b). *Garfinkel and Ethnomethodology*. Polity, Cambridge.
- Heritage, J. C. (1998). Harold garfinkel. In *Key Sociological Thinkers*. Macmillan Press.
- Herring, S. (1997). Humour that binds: Joking and coherence in internet relay chat. Paper presented at the International Society for Humor Studies conference. University of Oklahoma.
- Herring, S. C. (1999). Interactional coherence in cmc. *Journal of Computer Mediated Communication*, 4.
- Isaacs, E. and Clark, H. H. (1987). References in conversation between experts and novices. *Journal of Experimental Psychology*, 116:26–37.
- Jakobson, R. (1960). Closing statements: Linguistics and poetics. In A., S. T., editor, *Style in Language*. Wiley and MIT Press.
- Jefferson, G. (1987). On exposed and embedded correction in social interaction. In Button, G. and Lee, J. R. E., editors, *Talk and Social Organisation*, pages 86–100. Multilingual Matters, Philadelphia.

- Keenan, E. (1974). The universality of conversational postulates. In Shuy, R., editor, *Studies in Linguistic Variation*. Georgetown University Press.
- Koschmann, T. and LeBaron, C. (2003). Reconsidering common ground: Clark's contribution theory in the operating room. In *Paper presented at the European Computer-Supported Cooperative Work (ECSCW '03), Helsinki, Finland*.
- Krauss, R. M. and Fussell, S. R. (1996). Social psychological models of interpersonal communication. In Higgins, E. and Kruglanski, A., editors, *Social Psychology: Handbook of Basic Principles*, pages 655–701. Guildford Press, London.
- Krauss, R. M. and Weinheimer, S. (1966). Concurrent feedback, confirmation and the encoding of referents in verbal communication. *Journal of Personality and Social Psychology*, 4:343–346.
- Kroch, A. (1972). *The Semantics of Scope in English*. MIT.
- Lakoff, G. (1993). The contemporary theory of metaphor. In Ortony, A., editor, *Metaphor and Thought*. Cambridge University Press, Cambridge, 2nd edition.
- Larsson, S. (2005). Semantic plasticity through feedback and accommodation. Miscommunication Workshop, QMUL.
- Levinson, S. C. (1983). *Pragmatics*. Cambridge University Press, Cambridge.
- Levinson, S. C. (1992). Activity types and language. In *Talk at work*. Cambridge University Press, Cambridge.
- Lewis, D. (1969). *Convention: A philosophical study*. Harvard University Press.
- Linell, P. (1982). The written language bias in linguistics. In *Tema Kommunikation*. Linköping University.
- Locke, J. (1690). *Essay concerning Human Understanding*. Clarendon Press, Oxford.
- Macura, Z. and Ginzburg, J. (2006). Acquiring words across generations: introspectively or interactively? In *Proceedings of the 10th Workshop on the Semantics and Pragmatics of Dialogue*, University of Potsdam.

- McHoul, A. (1978). The organization of turns at formal talk in the classroom. *Language in Society*, 7:183–213.
- Metzing, C. and Brennan, S. (2003). When conceptual pacts are broken: Partner-specific effects on the comprehension of referring expressions. *Journal of Memory and Language*, 49:201–213.
- Mills, G. and Healey, P. G. T. (2006). Clarifying spatial descriptions: Local and global effects on semantic co-ordination. In *Proceedings of the 10th Workshop on the Semantics and Pragmatics of Dialogue*.
- O’Connell, D. and Kowal, S. (1995). Language use from a psychological perspective. In Radcliffe, F., editor, *Proceedings of the conference: Der Dialog-Begriff am Ende des 20. Jahrhunderts*, Berlin. Hegel-Institut.
- Pickering, M. J. and Garrod, S. (2004). Towards a mechanistic psychology of dialogue. *Behavioural and Brain Sciences*, 27(2):169–190.
- Pomerantz, A. (1984). Pursuing a response. In Atkinson, M. and Heritage, J., editors, *Structures of Social Action: Studies in Conversation Analysis*, pages 152–163. Cambridge University Press, Cambridge.
- Posner, M. I. and Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77:353–363.
- Purver, M., Healey, P. G. T., King, J., and Mills, G. (2003). Answering clarification questions. In *Proceedings of the 5th Workshop of the ACL SIG on Discourse and Dialogue (SIGdial03)*, Sapporo, Japan.
- Reddy, M. (1979). The conduit metaphor: a case of frame conflict in our language about language. In Ortony, A., editor, *Metaphor and Thought*. Cambridge University Press, London.
- Rosaldo, M. (1982). The things we do with words. Ilongot speech acts and speech act theory in philosophy. *Language in Society*, 11:203–237.
- Sacks, H. (1992). Spring 1970, 8 lectures. In Jefferson, J., editor, *Lectures on Conversation*. Blackwell, Oxford, UK.

- Sacks, H., Schegloff, E. A., and Jefferson, G. (1974). A simplest systematics for the organisation of turn taking in conversation. *Language*, 50:696–735.
- Sarangi, S. and Slembrouck, S. (1992). Non-cooperation in communication: a reassessment of gricean pragmatics. *Journal of Pragmatics*, 17:117–154.
- Saussure, F. (1974). *Course in General Linguistics : Trans W. Baskin*. Fontana, Glasgow.
- Schegloff, E. (1968). Sequencing in conversational openings. *American Anthropology*, 70:1075–1095.
- Schegloff, E. A. (1972). Notes on a conversational practice: Formulating place. In Sudnow, D., editor, *Studies in Social Interaction*. Free Press, Glencoe, IL.
- Schegloff, E. A. (1980). Preliminaries to preliminaries: ‘can i ask you a question?’. *Sociological Inquiry*, 50:104–0152.
- Schegloff, E. A. (1982). Discourse as an interactional achievement: Some uses of ‘uh huh’ and other things that come between sentences. In *Analyzing discourse: Text and talk*. Georgetown University Press.
- Schegloff, E. A. (1984). On some questions and ambiguities in conversation. In Atkinson, M. and Heritage, J. C., editors, *Structures of Social Action*. Cambridge University Press, Cambridge.
- Schegloff, E. A. (1987a). Analyzing single episodes of interaction: an exercise in conversation analysis. *Social Psychology Quarterly*, 50(2 Special Issue: Language and Social Interaction):101–114.
- Schegloff, E. A. (1987b). Some sources of misunderstanding in talk-in-interaction. *Linguistics*, 25(1):201–218.
- Schegloff, E. A. (1988). Presequences and indirection: Applying speech act theory to ordinary conversation. *Journal of Pragmatics*, 12:55–62.
- Schegloff, E. A. (1992). Repair after next turn: The last structurally provided defense of intersubjectivity in conversation. *American Journal of Sociology*, 97(5):1295–1345.
- Schegloff, E. A. (1997). Third turn repair. In Guy, G. R., Feagin, C., Schiffrin, D., and Baugh, J., editors, *Towards a Social Science of Language. Papers in honour of William Labov*, volume 2, pages 31–40. Benjamins, Amsterdam.

- Schegloff, E. A. (2000). When 'others' initiate repair. *Applied Linguistics*, 21(2):205–243.
- Schegloff, E. A., Jefferson, G., and Sacks, H. (1977). The preference for self-correction in the organization of repair in conversation. *Language*, 53:361–382.
- Schegloff, E. A. and Sacks, H. (1973). Opening up closings. *Semiotica*, 7:289–327.
- Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press, Cambridge, MA.
- Schlangen, D. (2004). Causes and strategies for requesting clarification in dialogue. In *Proceedings of the 5th SIGdial Workshop on Discourse and Dialogue*.
- Schober, M. F. and Clark, H. H. (1989). Understanding by addressees and overhearers. *Cognitive Psychology*, 21:211–232.
- Schutz, A. (1964). *Collected Papers*. Martinus Nijhoff.
- Schwartz, D. L. (1995). The emergence of abstract representations in dyad problem solving. *Journal of the Learning Sciences*, 4(4):321–354.
- Schwartz, D. L. (1999). The productive agency that drives collaborative learning. In Dillenbourg, P., editor, *Collaborative learning: Cognitive and computational approaches*, pages 197–218. Elsevier Science, New York.
- Schwartz, D. L. and Black, T. (1999). Inferences through imagined actions: knowing by simulated doing. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 25:116–136.
- Searle, J. R. (1969). *Speech Acts*. Cambridge University Press, Cambridge.
- Searle, J. R. (1972). What is a speech act? In *Language and Social Context*. Penguin, Harmondsworth.
- Searle, J. R. (1975). Indirect speech acts. In Cole, P. and Morgan, J., editors, *Syntax and Semantics. Speech acts*, volume 3. Academic Press.
- Searle, J. R. (1979). *Expression and Meaning : Studies in the Theory of Speech Acts*. Cambridge University Press.
- Searle, J. R. (1983). *Intentionality: an Essay in the Philosophy of Mind*. Cambridge University Press.

- Searle, J. R. (1990). Collective intentions and actions. In Cohen, J. M. P. R., Pollack, E., and Cohen, M., editors, *Intentions in Communication*, pages 401–416. MIT Press, Cambridge, Mass.
- Shannon, C. E. and Weaver, N. (1949). *The mathematical theory of communication*. Urbana: University of Illinois press.
- Sperber, D. and Wilson, D. (1986). *Relevance: Communication and Cognition*. Basil Blackwell Ltd.
- Sperber, D. and Wilson, D. (2002). Truthfulness and relevance. *Mind*, 111:583–632.
- Stahl, G. (2002). Shared meaning, common ground, group cognition. In *Computer support for collaborative learning: Foundations for a CSCL community: Proceedings of CSCL 2002*, pages 1–21. Lawrence Earlbaum Associates.
- Steels, L. (2003). The evolution of communication systems by adaptive agents. In *Adaptive Agents and Multi-Agents Systems*, pages 125–140. Springer Verlag.
- Steels, L. and Belpaeme, T. (2005). Coordinating perceptually grounded categories through language. a case study for colour. *Behavioural and Brain Sciences*, 28(4):469–489.
- Taylor, T. J. and Cameron, D. (1987). *Analysing Conversation: Rules and Units in the Structure of Talk*. Pergamon Press.
- Teichman, M. (2006). The germ of a sense. *Philosophy and Literature*, 2:567–579.
- Torode, B. (2002). Between uniqueness and universality: an ethnomethodological analysis of language games. In Grant, C. B., editor, *Rethinking Communicative Interaction*, page 209. Benjamins.
- Traum, D. (1999). Computational models of grounding in collaborative systems. In *Working notes of AAAI Fall Symposium on Psychological Models of Communication*, pages 124–131.
- Waismann, F. (1951). Verifiability. In Flew, A., editor, *Logic and Language*. Blackwell, London.
- Wilkes-Gibbs, D. and Clark, H. H. (1992). Coordinating beliefs in conversation. *Journal of Memory and Language*, 31:183–194.
- Wittgenstein, L. (1958a). *Blue & Brown Books*. Blackwell, Oxford.

Wittgenstein, L. (1958b). *Philosophical Investigations*. Blackwell, Oxford.

Appendix A

Maze game transcripts

A.1

(A.1) topgirl : blocked at?
 sat : thanks, can u get to ur goal now
 topgirl : yes
 topgirl : u dun?
 sat : same fing
 topgirl : ok
 topgirl : what u wana know?
 sat : listen, my goals at C3 R4
 sat : any gates open
 topgirl : ok mine is C6 R5
 topgirl : no
 topgirl : im one below ur goal
 sat : can u get to any of those
 topgirl : any of what?
 sat : any of C6 R2 or C5 R4
 sat : wheres ur nearest greys
 topgirl : ok i cant. u need to open my gate at C4 R2
 sat : where is ur nearest grey then, u need to guide me
 topgirl : ok um C4 R6
 sat : any other greys u got
 topgirl : or C3 R7
 sat : i cant get there, can u poss get to C6 R2 or C5 R4
 sat : check it properly
 sat : u shud b able to
 sat : odawise this maze is possible
 sat : sori, not poss
 topgirl : i cant get here
 topgirl : there
 topgirl : can u get to
 topgirl : wait wait
 topgirl : can u get to C3 R2
 sat : yes
 topgirl : go go
 sat : i jus did

A.2

Hei : and my switch is on third from top left
 Hei : no
 Rafa : guide my to your grey square
 Hei : where are you now again?
 Hei : ok
 Rafa : right side, third square from bottom
 Hei : third row from bottom still?
 Hei : ok
 Rafa : yes
 Hei : go left till end
 Hei : then go up (3rd square from top)
 Rafa : I can't, I'm blocked
 Hei : where?
 Hei : where is the switch?
 Hei : and where are you blocked?
 Rafa : MY swich is on the bottom row, second square from right
 Rafa : can you get there??
 Rafa : hei??
 Hei : I can't I'm blocked at 1 square from left
 Rafa : where is your switch??
 Hei : third from top
 Rafa : give me position, row, column
 Rafa : side?
 (A.2) Rafa :
 Rafa : man, say something
 Hei : my switch: column most left, on squire 3, from the top
 Hei : and me I'm on square 2, from left
 Hei : my gate is on squire 3, from left
 Rafa : I cannot reach your switch, I'm blocked
 Rafa : Do you have a switch on the right hand side??
 Hei : no :S
 Hei : only 1 switch
 Rafa : and whre is it? most left? second from bottom?
 Hei : where is your switch/gates?
 Hei : no third from bottom
 Hei : I mean square number 3 from bottom
 Rafa : can you go top??
 Rafa : go straight to, most left to
 Hei : yes top I can
 Hei : yeah
 Rafa : stay there
 Hei : I'm most left now
 Rafa : come on
 Rafa : top?
 Hei : now top
 Hei : oh
 Rafa : I'm in the goal now

A.3

Rafa : Go one box down
 Hei : ok
 Hei : done
 Hei : and 1 square left to the most left
 Rafa : wait.
 Rafa : Where are you now?
 Hei : 1 square left to middle
 Hei : I cant reach my 2 switch as there is a gate just above
 Rafa : do you have a switch on the top right side?
 Hei : top right just 2 square
 Hei : 2nd square
 Hei : from top right
 Rafa : now?
 Hei : ok the gate is open but raf
 Hei : where is your switches?
 Rafa : bottom row, middle box
 Hei : ok
 Hei : can you go back to your previous position?
 Rafa : i'm on goal now
 Hei : I'm blocked :D
 Hei : wait
 Hei : goal
 Rafa : I'm blocked now :(
 Hei : you have 2 switches?
 (A.3) Rafa : Can you go to most right, second reo from bottoom??
 Hei : here?
 Rafa : no
 Hei : hre?
 Rafa : no
 Hei : here?
 Rafa : second row, most right box
 Hei : here?
 Rafa : done
 Hei : :D
 Rafa : Where are you?
 Hei : I'm in middle
 Rafa : top or bottom?
 Rafa : no bottom but middle!! Third row from bottom
 Hei : where are you?
 Rafa : second reow from bottom, left most
 Hei : can you go : down, right?
 Rafa : yoes
 Hei : oh it doesnt does anything to my gate
 Rafa : second row from bottom
 Rafa : so, give me your position
 Hei : ok
 Rafa : I am in 2,6
 Rafa : where are you??
 Hei : 4, 4

A.4

(A.4) Hei : for the row you count from the top?
 Hei : ok
 Hei : yeah ok I see
 Rafa : ok, go to 6,5. Can you get there?
 Hei : sorry
 Hei : 4,4
 Rafa : yes
 Hei : wait
 Hei : there?
 Hei : I'm in 6,5
 Rafa : no??yes?? what!!
 Hei : ok
 Hei : yes I'm in it
 Hei : :D
 Rafa : where are you??
 Hei : 6,5
 Hei : oh no
 Hei : sorry
 Hei : 6,5
 Hei : 6,4
 Hei : I have a gate in...
 Rafa : go to 6,5
 Hei : 4,4 : gate blocking me
 Rafa : ok, do you have a switch on the bottom 6,6
 Rafa : give a position of one switch
 Hei : gate below 4,4
 Hei : 6.4
 Hei : and 3.7
 Hei : ok I'm in my goal, which is in 3,6
 Rafa : I'm in 3,7
 Rafa : My goal is in 6,3 but I need you to activate a switch in
 Hei : I'm now and 6,5 on your switch I think
 Rafa : I'm in goal now
 Hei : ok
 Hei : :D
 Rafa : give me position
 Hei : 4.5
 Rafa : I'm in 5,4
 Hei : go one left plz
 Rafa : can you go to 7,7
 Hei : no gate rightside
 Rafa : OK
 Rafa : I got a gate on my left
 Hei : my switches: 3.7 6.4 7,7
 Hei : can you go one of them?
 Rafa : can you go to 3,3

A.5

Hei : ok
 Rafa : can you reach any f them
 Hei : yeah one
 Hei : 7,3
 Rafa : ok
 Rafa : now, where are yours?
 Hei : I have only one switch on 1,7
 Rafa : Ok, I,m in 2,1 and my switches are in 4,3 and 7,4
 Rafa : can you reach any switch?
 Hei : my swtiches: 4.3 and 7,4
 Hei : erm
 Hei : one switch
 Hei : goal 2.1
 Hei : I'm in 5, 5 and my switches are 4,1 and 3, 5
 Rafa : I'm on 5,2 and my switches are in 3,2 and 2,5 and 6,6
 Hei : ok I move
 Hei : on
 Hei : 6.6
 Hei : done
 Hei : ?
 Rafa : nop
 Rafa : now
 (A.5) Hei : 2,5 now
 Hei : I'm on 2,5
 Rafa : open now?
 Hei : goal now : 5,2
 Rafa : done
 Hei : :D
 Hei : 7,6 now
 Hei : done anything?
 Rafa : yes
 Rafa :
 Hei : I'm on 3,5 I cant move upward
 Hei : blocked by gate
 Rafa : I cannot move now since I,m blocked in the goal box
 Hei : switches are: 1,3 1,7 5,6, 7,2
 Hei : ok
 Rafa : I need you to open first
 Rafa : now, go to your goal
 Hei : done
 Rafa : where is your goal?
 Hei : goal 1,7
 Hei : 1,6
 Rafa : can go to 1,3
 Hei : done

A.6

```

linuxblackhat : u look from the left....the third column and the lower most block
linuxblackhat : move to that block
linuxblackhat : my gate will open then
linuxblackhat : didnt u understand me ?
Ebby :         actually that is block for me
linuxblackhat : ok
linuxblackhat : then look from right
Ebby :         I can't go there
linuxblackhat : the third column
Ebby :         ok let me tell you now
linuxblackhat : and the 2nd block from the top
linuxblackhat : from right...3rd column and the second block from the top
linuxblackhat : go to tht block
Ebby :         open now
linuxblackhat : no
linuxblackhat : from the right ..its 3rd column
linuxblackhat : and from top its the 2nd block
Ebby :         I am in 2nd block on 3 from right
linuxblackhat : still not
Ebby :         ok
linuxblackhat : how many blocks do u see in that column
Ebby :         7
linuxblackhat : thats the fourth colum
(A.6) Ebby :         i m right
linuxblackhat : yes its open now
linuxblackhat : where are u
linuxblackhat : hmm
linuxblackhat : do one thing
linuxblackhat : again go to the same blocok
linuxblackhat : where u went b4
linuxblackhat : so that my gate opens
linuxblackhat : from right its the 3rd column and 2nd block from the top
Ebby :         can u come on 4 column and 6 block
linuxblackhat : my way is blocked
linuxblackhat : thts y i m saying
linuxblackhat : that go to the same block u went b4
linuxblackhat : so that my dooor opens
Ebby :         I m on cross
linuxblackhat : is ur door open now
Ebby :         yes
linuxblackhat : bt now my door is locked for the goal
linuxblackhat : i m on the goal also
linuxblackhat : are u on the goa
linuxblackhat : goal ?
Ebby :         ok stay there
Ebby :         i m comming
linuxblackhat : (Y)

```

A.7

linuxblackhat : now do one thing
 Ebby : what
 linuxblackhat : can u go to 2nd column from right and the 2nd block
 linuxblackhat : from top
 linuxblackhat : ok u can go to ur gaol now ?
 Ebby : door for me is closed
 linuxblackhat : how many grey blocks do u have ?
 Ebby : 1
 linuxblackhat :
 linuxblackhat : ok
 Ebby : which is on 3 col and 2 blo from top on right side
 linuxblackhat : n where is it
 linuxblackhat : its open now ?
 Ebby : i m on cross
 Ebby : yes
 linuxblackhat : bt i can t go to crooss now
 linuxblackhat : can u go to the
 Ebby : which
 linuxblackhat : 3rd column from right and the last bottom block
 Ebby : yes
 Ebby : congrates
 linuxblackhat : ok where are u now ?
 linuxblackhat : ?
 Ebby : i m on 4 col 3 blo from top right
 (A.7) Ebby : i have 3 gray areas
 linuxblackhat : and where is ur cross
 linuxblackhat : ?
 Ebby : 2 col and top blo from right
 linuxblackhat : ok now
 linuxblackhat : can u go to
 Ebby : wher
 Ebby : where
 linuxblackhat : the first column and the bottom most blocok from right
 Ebby : no
 linuxblackhat : can u go to
 linuxblackhat : 3rd column from left and the top most block
 linuxblackhat : ?
 Ebby : is it open
 linuxblackhat : can u go to
 Ebby : where
 linuxblackhat : the first column bottom block from right
 Ebby : i m there
 Ebby : already
 linuxblackhat : now i m therw
 linuxblackhat : ur door is open right
 Ebby : i m on cross
 linuxblackhat : bt mine is closed now
 Ebby : where are you now
 linuxblackhat : note down..i have 2 grey area... one is from 3rd colum from left block 1

A.8

Chris : 1,1 1,2 1,3 1,4 1,5
 Chris : 2.1 2.2 2.3 2.4 2.5
 Chris : 3.1 3.2 3.3 3.4 3.5
 Chris : you understand this matrix?
 lgm : ok
 Chris : we say it like this
 Chris : Now I am at 3,7
 Chris : I am at the switch
 Chris : any effect on you?
 lgm : yes
 Chris : my switch is 4,5 and 6,4
 Chris : lgm? any problem?
 (A.8) Chris : the unblock is still unblocked...
 lgm : the gate is blocked again
 Chris : now?
 lgm : I can't go to your swith
 Chris : why?
 lgm : 5,6
 Chris : say our swi at fisrt time
 Chris : so your swi?
 Chris : I cannot go there either!
 lgm : 2, 6 or 4,1 or 6,5
 Chris : i ma in 2,6
 Chris : can you go to 2, 5
 Chris : i am in goal

A.9

betina : ok sweetie, maybe we cud use rows and colomns to work out location
 betina : wht
 Katy : look at it as a grid ok
 betina : yep
 Katy : colums go down
 Katy : rows are across
 betina : where r u
 betina : yep
 Katy : and the fourth box down
 Katy : thats one of my grey areas
 betina : yep, b4 that 4 me theres a gate
 betina : thas b4 my x
 Katy : ok tell me where one of your grey ares is
 betina : my dot is the 1 col n 3 row
 Katy : i dont need to know where your dot is i need to know your grey area
 betina : 1 col n 5 row
 betina : i cant get to grey coz theres a gate in my way grrrrrr
 Katy : ok i cant get there either
 Katy : ok babes
 Katy : there am on it
 betina : yes second row down n 5 col along
 betina : im on my x
 betina : right wheres urs
 betina : grey area
 (A.9) Katy : ok take your dot to colum five 4th down
 betina : say it again
 betina : cant readv screen
 Katy : fifth colum 4th row
 Katy : thats one of my grey areas the fifth colum fourth row
 Katy : get your dot on there
 betina : cant the gates blockin it go bk to my grey area
 betina : again
 betina : ok im thr
 Katy : am there
 Katy : on my X
 Katy : are u
 Katy : ta bird
 betina : no coz i had to come off 2 help u
 betina : oh this is hard!!!!!!!!!!!!!!
 Katy : ok wher is one of your grey area
 betina : u were just on it
 betina : right im on x
 Katy : me too
 Katy : ok
 betina : ok
 Katy : tell me one of your grey araes
 betina : im on 5 row n 6 col
 betina : 3 row n 7 col

A.10

(A.10) betina : im on row 4
 Katy : it hasnt let me through
 Katy : its colum 7
 Katy : the end colum
 Katy : and four boxes down
 betina : well im on i cant get there
 betina : another grey
 Katy : ok u cant get there
 Katy : my other grey is colum 4
 Katy : row 3
 betina : u thru
 Katy : yes am thru
 betina : go bk u just opened my gate
 Katy : your grey is ?
 Katy : i cant go back go back to my grey
 betina : 4,3 or 7,4
 Katy : either yes
 betina : on it
 Katy : am not it blocked me
 Katy : do u just have two grey
 betina : ok another grey yes 2
 betina : do u have another grey
 Katy : i have two mine are colum 4 row 3
 betina : and
 Katy : and colum 7 row 4
 betina : ok
 Katy : yeah
 Katy : u?
 Katy : well donw
 Katy : my grey is colum 1 row 4
 Katy : and colum 5 row 1
 Katy : where are yours
 betina : u thru
 Katy : yes
 Katy : where are ur greys
 betina : mine r 1,2 3, 7 n 7,6
 Katy : betina where are your greys?
 Katy : answer me women
 betina : iv said
 Katy : i haven got them
 betina : 1, 2 3, 7 n 7,6
 Katy : ok let me through again
 betina : where
 betina : there
 Katy : ok u thru
 betina : no
 betina : yes on x
 Katy : yes well donw

A.11

(A.11)

pabz : open?
slink : nop
pabz : hmm
slink : ok u see the starting from the left the 3rd vertical thing yeh
slink : the bottom box on that
slink : u opened it and closed it lol
pabz : that?
slink : nop
pabz : now?
slink : nah
slink : its the last box on the third column
slink : ok
pabz : u through?
slink : its box 11
pabz : wtf?
pabz : u see where it says maze at the top left of all
slink : yeh
pabz : down that column u see 3 boxes ye
slink : yeh
pabz : then 3
pabz : then 5
slink : yeh
pabz : then 7
slink : yeh
slink : its the one with 5
slink : the bottom one
pabz : which column of that
pabz : k
pabz : now
slink : didnt open
pabz : the box at the top is square 1
pabz : which one?
pabz : 2? 3?
pabz : ignore the second
slink : yeh
pabz : where ur gate is
slink : yeh
pabz : two boxes down is mine
slink : therr?
pabz : through
pabz : ye
pabz : now
pabz : it shud b open
pabz : urs
slink : yeh
pabz : got ox
slink : i am on the X
pabz : finally lol

A.12

(A.12) pabz : or cant u?
slink : k
slink : has it opened?
pabz : there
pabz : X
pabz : first column, box 3
pabz : k
slink : yeh
slink : first column box 1
pabz : wher are ur switches?
pabz : im there
slink : ok i can get on the X now
slink : can u?
pabz : nop
pabz : column 5
pabz : first box
slink : yeh?
pabz : X
pabz : column 3 box 3
slink : im hter
pabz : column 6 box 2
pabz : im on x
pabz : those are my 2 switches
slink : column 6
slink : box 3
slink : k on X
pabz : X
pabz : start by saying ur switches
slink : column 3 box 1
pabz : column 5 last box
pabz : column 7, only box
slink : cnt get ther
slink : to either
pabz : sorry
slink : cnt get to that 1
pabz : column 4 fist box?
pabz : now urs?
slink : wers ur goal?
pabz : column 5 box 2
slink : can u get to column 6 last box
slink : can u get to column 2 box 2
pabz : ye
pabz : there
pabz : now open mine again
slink : x
slink : k
slink : wer is urs
pabz : column 3 box 4

A.13

chap : so my grey square starting from tbottom left square is 7 up and 3 across
 pat : to unlock the gate
 chap : look i can't get there so tell me where u are
 pat : I am on the 3rd sqare on the horizontal row of 6
 pat : where r u?
 pat : G where r u?
 chap : ok i am on the 5th square on the 5th row
 pat : where is ur gate
 chap : go back
 chap : just opened
 chap : start
 pat : go 1st
 pat : f5
 chap : we are gettin good
 pat : bout u?
 pat : at goal
 pat : can u get to goal?
 chap : need u to get to aither d4, b5, f6
 pat : at d4
 pat : u at goal yet?
 chap : i meant d1
 pat : cant go there
 pat : i'll go to b5
 (A.13) chap : none?
 chap : i'm at goal
 pat : start
 chap : a3,b5,f3
 pat : at f3
 chap : at goal
 pat : d5 or f6
 chap : so slow
 pat : get t oat goal
 chap : well done
 pat : u 1st
 chap : a6
 chap : no c6
 pat : no can do
 chap : gimme urs fool
 pat : a3 or d6
 chap : stop crashing
 chap : opened
 pat : at goal
 chap : help me got co ordnates
 pat : can u get to goal?
 chap : no u need to get to c6
 pat : where r u?
 pat : at c6

A.14

(A.14) chap : fine
chap : need a fag
pat : at goal yet?
chap : yes
pat : me too
pat : ladies 1st
chap : go on then
pat : u
pat : cant go any where
chap : d7 and e4
chap : bitch
pat : get to d7 or e4
pat : can u get to either of d7 or e4
pat : at goal
chap : opened up for me go back
pat : bout u
chap : found it
chap : not the goal
pat : now im blocked at goal
chap : opened
pat : im out
chap : u need to let me g first to goal
chap : the g was go
pat : where u want to go
chap : get to e4
chap : fin
chap : finito for me
chap : u can get there
pat : where is your goal
chap : at goal
pat : my goal is g2
chap : b7 is goal
pat : can you get ther now
chap : yes

A.15

(A.15) kaz : the switch i mean
alex : which?
kaz : no the gate hasn't opened
kaz : try the other switch
kaz : ok im on a switch
alex : i do not have other
alex : they r on the bottom
kaz : has ur gate opened
kaz : ?
kaz : stay where u r
alex : my nt
kaz : uh?
alex : u have to go first down and open mine
alex : gate
kaz : no but my switch opens ur gate
alex : your not mine
alex : stay
kaz : u on the target?
alex : im in
alex : r u?
kaz : ok can u get on a switch to open my gate now?
alex : which one?
kaz : no the other switch
alex : tell me where i shoud go
kaz : well where r u?
kaz : and stay still
alex : 6 row 3 sq - my target
kaz : alex where r u?
kaz : go up 1
alex : ok
kaz : and 2 left
kaz : stay
kaz : im there
alex : u on the target
kaz : u can go back to yours?
kaz : yes on target
alex : yes :))
kaz : finally!
alex : fuck :))
kaz :
kaz : ok
kaz : first column, 2nd from bottom
kaz : is where i am
alex : i am
kaz : go along to the left
kaz : i.e.
alex : ok
kaz : first column

A.16

(A.16) girl1 : did u get to your goal
 girl1 : no
 bakingwithsylvania: k, hold on
 girl1 : 2nd row from top
 bakingwithsylvania: there?
 girl1 : there
 bakingwithsylvania: k.
 girl1 : ok i'm at my goal
 bakingwithsylvania: ok. now
 girl1 : do we both need to be there at the same time?
 bakingwithsylvania: yeh, i thnk so :(
 bakingwithsylvania: ok. I'm currently 4th from the top, second from the left.
 girl1 : ok do you have a switch on the top right part
 bakingwithsylvania: I have a gate right under me.
 bakingwithsylvania: yes, very top box
 girl1 : good?
 bakingwithsylvania: awesome
 bakingwithsylvania: in goal
 bakingwithsylvania: ok
 bakingwithsylvania: YAY
 bakingwithsylvania: ok, my switch is in the second from top row on the left
 bakingwithsylvania: where's yours?
 bakingwithsylvania: I have gates between the top segment and the wide segment
 bakingwithsylvania: goal is second from leftmost box
 bakingwithsylvania: do you have any switches int he part below the top three rows?
 girl1 : I am 2nd from top
 girl1 : yes
 bakingwithsylvania: your goal or your swtich?
 girl1 : where are you
 bakingwithsylvania: in my goal, second from leftmost box
 bakingwithsylvania: but I can move pretty much anywhere except the top three rows
 girl1 : ok my switch is the leftmost box
 bakingwithsylvania: so do you have a switch in the bottom part of the maze?
 bakingwithsylvania: ok, I have a gate there too
 girl1 : where is the switch
 bakingwithsylvania: second from top on the left
 bakingwithsylvania: k, in goal. can you get back to yours ok?
 girl1 : no
 bakingwithsylvania: where're your gates?
 girl1 : my gates are on the rightmost column
 bakingwithsylvania: that's fine
 bakingwithsylvania: ok
 bakingwithsylvania: hmm
 girl1 : between the 3rd fand 4th row
 bakingwithsylvania: wait, 3rd and 4th row from the top? :-/
 bakingwithsylvania: in the rightmost?
 girl1 : if i get in the bottom 3 rows i am good
 girl1 : i have goal and switches there

A.17

bakingwithsylvania: yours?
 girl1 : switches: 1) top row, 3rd from left
 bakingwithsylvania: k, can't get there yet...
 girl1 : 2) 3rd from right, 4th from bottom
 bakingwithsylvania: open?
 girl1 : no
 girl1 : counting the space
 bakingwithsylvania: did that open it?
 girl1 : no
 bakingwithsylvania: there?
 girl1 : yay
 bakingwithsylvania: ok, in goal. you good?
 bakingwithsylvania: ok. one switch bottom row, third from left
 bakingwithsylvania: and one 4th from bottom row, 2nd from right
 bakingwithsylvania: awesome, I'm good
 bakingwithsylvania: where are yours?
 girl1 : ok one switch, 2nd from right, 3rd from bottom
 bakingwithsylvania: any in the bottom three rows?
 bakingwithsylvania: good?
 girl1 : go to goal
 bakingwithsylvania: k
 bakingwithsylvania: one 2nd from right, 2nd row down
 girl1 : switches: bottom row, 3rd from left
 (A.17) bakingwithsylvania: and one 3rd from right, 4th row down
 bakingwithsylvania: you?
 bakingwithsylvania: there's the switch where yours is, 3rd from top, 4th from left
 girl1 : um
 bakingwithsylvania: and the bottom one, 4th from top, rightmost
 bakingwithsylvania: ok, go to goal
 girl1 : open?
 bakingwithsylvania: yup
 girl1 : my goal is rightmost, top
 bakingwithsylvania: by your goal, right?
 girl1 : where are your switches
 girl1 : yes, a gate cuts off my goal
 girl1 : so i need to let you through first
 bakingwithsylvania: hmm.
 bakingwithsylvania: ok.
 girl1 : where are your switches
 girl1 : no
 girl1 : yes
 bakingwithsylvania: k :)
 bakingwithsylvania: oh! sorry
 bakingwithsylvania: where's the bottom one again?
 girl1 : 2nd from bottm, 3rd from left
 bakingwithsylvania: there?
 girl1 : go to goal