

Open Research Online

The Open University's repository of research publications and other research outputs

Non-cooperation in dialogue

Conference or Workshop Item

How to cite:

Plüss, Brian (2010). Non-cooperation in dialogue. In: Proceedings of the Student Research Workshop at the 48th Annual Meeting of the Association for Computational Linguistics, 11-16 Jul 2010, Uppsala, Sweden.

For guidance on citations see [FAQs](#).

© 2010 Association for Computational Linguistics (ACL)

Version: Accepted Manuscript

Link(s) to article on publisher's website:
http://acl2010.org/program_mainconf.html#s53

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's [data policy](#) on reuse of materials please consult the policies page.

oro.open.ac.uk

Non-Cooperation in Dialogue

Brian Plüss

Centre for Research in Computing

The Open University

Milton Keynes, UK

b.pluss@open.ac.uk

Abstract

This paper presents ongoing research on computational models for non-cooperative dialogue. We start by analysing different levels of cooperation in conversation. Then, inspired by findings from an empirical study, we propose a technique for measuring non-cooperation in political interviews. Finally, we describe a research programme towards obtaining a suitable model and discuss previous accounts for conflictive dialogue, identifying the differences with our work.

1 Introduction

Most approaches to modeling conversation are based on a strong notion of cooperation between the dialogue participants (DPs). Traditional models using intentions (Cohen and Levesque, 1991), dialogue games (Power, 1979), shared plans (Grosz and Sidner, 1990) or collaborative problem-solving (Blaylock and Allen, 2005) explain dialogue situations in which DPs recognise each other's intentions and, at least to a certain extent, accept each other's goals when deciding on their actions. These assumptions are theoretically grounded, as most work in linguistics has considered situations in which DPs share a common goal and cooperate to achieve it by means of conversation (Grice, 1975; Clark and Schaefer, 1989). They are also practically sound: dialogue models are usually implemented in the form of dialogue systems, built for the purpose of providing a service to their users (e.g., TRAINS (Allen and Schubert, 1991)). In this scenario, failure to cooperate, either on the side of the system or of the user, is against the premises on which the system is conceived and used.

In everyday conversation, however, a great many situations escape the arguments above. Con-

sider the following example¹:

- (1) PAXMAN [1]: (interrupting) Did you threaten to overrule him?
HOWARD [2]: I, I, was not entitled to *instruct* Derek Lewis, and I did not instruct him.
PAXMAN [3]: Did you threaten to overrule him?
HOWARD [4]: The truth of the matter is that Mr. Marriott was not suspended. I...
PAXMAN [5]: (overlapping) Did you threaten to overrule him?
HOWARD [6]: ... did not overrule Derek Lewis.
PAXMAN [7]: Did you *threaten* to overrule him?
HOWARD [8]: I took advice on what I could or could not do...
PAXMAN [9]: (overlapping) Did you threaten to overrule him, Mr. Howard?
HOWARD[10]: ... and I acted scrupulously in accordance with that advice, I did *not* overrule Derek Lewis...
PAXMAN[11]: (overlapping) Did you threaten to overrule him?
HOWARD[12]: ... Mr. Marriott was *not* suspended.
PAXMAN[13]: Did you threaten to overrule him?
HOWARD[14]: (pauses) I have accounted for my decision to dismiss Derek Lewis...
PAXMAN[15]: (overlapping) Did you threaten to overrule him?
HOWARD[16]: ... in great detail, before the House of Commons.
PAXMAN[17]: I note that you're not answering the question of whether you *threatened* to overrule him.

(*Newsnight*, BBC, 1997)

We take it for granted that, at some level, Paxman and Howard are sharing a goal, for otherwise they would not be having an interview. Still, the exchange is clearly conflictive, to the point that their behaviour compromises the flow of the conversation.

Heritage (1998) analyses the distinctive roles of DPs in news interviews:

¹BBC presenter Jeremy Paxman questions former UK Home Secretary Michael Howard with respect to a meeting in 1995 between Howard and the head of the Prison Service, Derek Lewis, about the dismissal of the governor of Parkhurst Prison, John Marriott, due to repeated security failures. The case was given considerable attention in the media, as a result of accusations by Lewis that Howard had instructed him, thus exceeding the powers of his office.

“the participants -IRs [=interviewers] and IEs [=interviewees]- exclude themselves from a wide variety of actions that they are normally free to do in the give and take of ordinary conversation. If IRs restrict themselves to asking questions, then they cannot - at least overtly - express opinions, or argue with, debate or criticize the interviewees’ positions nor, conversely, agree with, support or defend them. Correspondingly, if IEs restrict themselves to answers (or responses) to questions, then they cannot ask questions (of IRs or other IEs), nor make unsolicited comments on previous remarks, initiate changes of topic, or divert the discussion into criticisms of the IR or the broadcasting organization.”

(Heritage, 1998, p.8)

Now, consider the fragment below²:

- (2) PAXMAN[1]: Can you clear up whether or not you did threaten to overrule Derek Lewis when you were Home Secretary?
 HOWARD[2]: Oh, come on, Jeremy, you are really going to go back over that again? As...
 PAXMAN[3]: (overlapping) You’ve had seven years to think about it!
 HOWARD[4]: (overlapping)...as, as it happens, I didn’t. Are you satisfied now?
 PAXMAN[5]: Thank you. Why didn’t you say that at the time?
 HOWARD[6]: I, well, we’ve been over this many, many times. I, I, I knew that everyone was crawling over every syllable I said about that, and I wanted to check very carefully what I said before answering your question.

(*Newsnight*, BBC, 2004)

On this occasion, Howard provides an answer almost immediately and the flow of the conversation contrasts noticeably with that in (1). The investigation reported in this article aims at shedding light on the nature of non-cooperation in dialogue, by capturing the intuitions that allow us to differentiate between both conversations in terms of participant behaviour.

Dialogue games supporters could say that there is a game that describes the interaction in the first example. While this might be true, such an approach would force us, in the limit, to define one game for each possible conversation that would not fit a certain standard. Walton and Krabbe (1995) attempt a game-based approach in their study of natural argumentation. They claim that a rigorous model of conversational interaction is useful, but accept that most of the huge variety of everyday conversation escapes it. Dialogue games are based on strict rules that capture typical dialogue situations while leaving out considerable detail. As example (1) shows, DPs behaviour can

²This exchange took place seven years after (1), when public awareness of the 1995 affair had dissipated.

divert from the typical case in unexpected ways, falling outside the characterisation³.

Nevertheless, the rules and patterns captured by game models are useful, as they describe the expected behaviour of the DPs under a certain conversational scenario. In our research, we aim at reconciling two worlds, using the insights from dialogue games to provide a description of expected behaviour in the form of social obligations, but looking at naturally occurring cases that deviate from the norm. This, in turn, calls for a technique to measure non-cooperation in dialogue and in this paper we provide one that is theoretically sound and supported by empirical evidence.

The following section discusses levels of cooperation in dialogue; Section 3 presents an empirical study and a practical measure of non-cooperation in political interviews; in Section 4 we discuss related work, our working hypothesis and a methodology; and Section 5 has the conclusions.

2 Linguistic and Non-Linguistic Cooperation

Cooperation in dialogue can happen at different levels. In most cases, conversation supports a social activity that constrains the behaviour acceptable or expected from the participants. In addition, conversational behaviour determines how cooperatively participants engage in a social activity. However, cooperation at the conversational level does not necessarily translate to the social level. Consider, for instance, a witness under interrogation in a U.S. trial refusing to answer a question by appealing to the Fifth Amendment of the Constitution⁴. Such behaviour will be accepted in the conversational setting as established by law, although it is not cooperative in relation with the goals of the trial. Non-cooperation at the conversational level, on the other hand, usually results in lack of cooperation at the social level. Take as an example, the same witness remaining silent, rather than answering or appealing to the Fifth Amendment.

To illustrate further, consider a fictional alternative to (1), where Howard replies by saying “I will not answer that question, as it is not relevant to whether I exceeded the powers of my office”.

³Consider, for instance, Ginzburg’s QUD model (Ginzburg, 1996) when applied to dialogue (1), in which Howard repeatedly fails to either accept or reject Paxman’s question.

⁴“No person shall (. . .) be compelled in any criminal case to be a witness against himself”.

This is not cooperative for the interview, but it is so at the linguistic level. It would help in preserving the flow of the conversation, e.g., by triggering a sub-dialogue to solve the disagreement.

The distinction between linguistic and non-linguistic (also called task-related, high-level or social) cooperation has been addressed before. Attardo (1997) revisits Gricean pragmatics, relating non-linguistic cooperation to participants' behaviour towards realising task-related goals, and linguistic cooperation to assumptions on their respective behaviour in order to encode and decode intended meaning. From a computational perspective, Bunt (1994) relies on a similar distinction for defining dialogue acts. Also, Traum and Allen (1994) introduce discourse obligations as an alternative to joint intentions and shared plans, to allow for models of dialogues in which participants do not share the same high-level goals and where behaviour is also determined by "a sense of obligation to behave within limits set by the society" (Traum and Allen, 1994, p.2).

Walton and Krabbe (1995) proposed a typology of dialogue based on the initial situation triggering the exchange and participants' shared aims and individual goals. Based on their work, Reed and Long (1997) distinguish cases where participants follow a common set of dialogue rules and stay within a mutually acknowledged framework from a stronger notion in which their individual goals are in the same direction. Borrowing from the latter, in the rest of the paper, we will speak of *collaboration* when DPs share the same task-level goals, and use *cooperation* when participants follow the conversational obligations imposed by the social activity (i.e., linguistic cooperation as discussed above). We will not deal with collaboration here, though, as our focus is on non-cooperation.

3 An Empirical Study

In this section, we describe an empirical pilot study aimed at identifying a set of features that distinguish cooperative from non-cooperative conversational behaviour and at establishing a suitable domain in which to focus our work.

3.1 The Corpus

We collected the transcripts of 10 adversarial dialogues: 4 political interviews, 2 entertainment interviews, 1 parliamentary inquiry, 1 courtroom confrontation, 1 courtroom interrogation and 1

dispute. The corpus includes 2 collaborative political interviews for result comparison and is nearly 14,500 words long⁵.

In a first analysis, we identified those surface features that characterised each conversation as conflictive: e.g., interruptions, short turns, unfinished adjacency pairs, verbatim repetition. Next, looking for a better understanding, we performed an in-depth case study of one of the examples, approaching the analysis from different angles.

By studying, e.g., the observance of turn-taking rules, the implicatures of the participants and, more extensively, how the case fitted within the normative framework proposed by Walton and Krabbe (1995), we were able to better identify the nature of non-cooperative features present in the dialogue and establish a formalisable framework for approaching non-cooperative dialogue.

As for the domain, the wealth of interesting conversational situations that arise in political interviews make a suitable context for this research. In the English-speaking world, journalists are well-known for their incisive approach to public servants. At the same time, politicians are usually well trained to deliver a set of key messages when speaking in public, and to avoid issues unfavorable to their image. We will only consider naturally occurring (i.e. non-scripted) two-party interviews.

3.2 Degrees of Non-Cooperation

Based on the analysis described above, we propose a technique for measuring non-cooperation in political interviews using a set of non-cooperative features (NCFs). The number of occurrences of these features will determine the degree of non-cooperation (DNC) of an exchange.

We grouped NCFs following three aspects of conversation: turn-taking, grounding and speech acts (see Table 1 for a complete list).

Turn-taking rules (Sacks et al., 1974) establish that speakers make their contributions at adequate places and in particular ways. Interlocutors in a political interview are expected to respect transition-relevance places, openings and closings according to social conventions. Failing to do so (e.g., by interrupting each other) constitutes a non-cooperative feature.

Grounding (Clark and Schaefer, 1989) refers to participants' acknowledgement of each other's

⁵These resources are available at <http://www.open.ac.uk/blogs/brianpluss/pilot-study/>.

Turn-Taking	For both speakers: <ul style="list-style-type: none"> • interrupting • overlapping • ending the exchange abruptly
Grounding	Interviewer fails to either: <ul style="list-style-type: none"> • ask next relevant question • move to next topical issue • state irrelevance of answer Interviewee fails to either: <ul style="list-style-type: none"> • give relevant answer • reject question
Speech Acts	Interviewer either: <ul style="list-style-type: none"> • expresses personal opinion • argues, debates with or criticises interviewee’s position subjectively • agrees with, supports or defends interviewee’s position subjectively Interviewee either: <ul style="list-style-type: none"> • asks (non-CR) question • makes irrelevant comment • initiates change of topic • criticises interviewer

Table 1: NCFs for political interviews

contributions by providing evidence of understanding (e.g, continued attention, relevant next contribution). In political interviews a question is acknowledged by rejecting it or by providing a direct answer. Likewise, answers are acknowledged by rejecting their relevance, by asking a next relevant question or by moving on to a new topical issue. Failing to provide sufficient evidence of understanding is also a non-cooperative feature.

Speech Act theory (Searle, 1979) classifies utterances according to their associated force and propositional content. Going back to Heritage’s comment, in a political interview participants can fail to restrict their speech acts to the force and content expected for their role. Non-cooperative features related to speech acts include the interviewer expressing a personal opinion or criticising subjectively the interviewee’s positions and the interviewee asking questions (except for clarification requests) or making irrelevant comments.

We define the degree of non-cooperation (DNC) of a dialogue as the proportion of utterances with one of more occurrences of these non-cooperative features⁶. Furthermore, the DNC could be thus computed for the whole conversation and also for each participant, by counting only occurrences of features and utterances from each DP.

As an example, consider an extended fragment

⁶At this stage, all NCFs are weighted equally. This is a simplifying assumption we will remove in the future so that, e.g., an interviewee attempting a change of topic has a stronger impact on the DNC than, say, one interrupting.

of (1) annotated with non-cooperative features (**O**: overlap; **GF**: grounding failure; **UC**: unsolicited comment; **I**: interruption; **TC**: topic change):

- (3) P [11] : Uir.1 (overlapping) Did you threaten to **O**
overrule him?
H[12] : Uie.1 ... Mr. Marriot was *not* suspended. **GF**
P [13] : Uir.2 Did you threaten to overrule him? **GF**
H[14] : Uie.2 (pauses) I have accounted for my decision to dismiss Derek Lewis. . .
P [15] : Uir.3 (overlapping) Did you threaten to **O**
overrule him?
H[16] : Uie.2 ... in great detail before the House of **UC**
Commons.
P [17] : Uir.4 I note that you’re not answering the
question whether you *threatened* to
overrule him.
H[18] : Uie.3 Well, the important aspect of this **GF**
which it’s very clear to bear in
mind. . .
P [19] : Uir.5 (interrupting) I’m sorry, I’m going to **I**
be frightfully rude but. . .
H[20] : Uie.4 Yes, you can. . .
P [21] : Uir.6 (overlapping) I’m sorry. . . **O**
H[22] : Uie.4 (overlapping) ...you can put the **O**
question and I will give you, I will
give you an answer.
P [23] : Uir.7 ...it’s a straight yes-or-no question
and a straight yes-or-no answer:
Uir.8 did you threaten to overrule him?
H[24] : Uie.5 I discussed the matter with Derek
Lewis.
Uie.6 I gave him the benefit of my opinion.
Uie.7 I gave him the benefit of my opin- **UC**
ion in strong language, but I did not
instruct him because I was not, er,
entitled to instruct him.
Uie.8 I was entitled to express my opinion **UC**
and that is what I did.
P [25] : Uir.9 With respect, that is not answering
the question of whether you threat-
ened to overrule him.
H[26] : Uie.9 It’s dealing with the relevant point **TC**
which was what I was entitled to do
and what I was not entitled to do,
Uie.10 and I have dealt with this in detail **UC**
before the House of Commons and
before the select committee.

Table 2 summarises non-cooperative features, utterances and the degree of non-cooperation for each participant and for the whole fragment.

	P (ir)	H (ie)	Fragment
Interruptions	1	0	1
Overlaps	3	1	4
Grounding Failure	1	2	3
Unsolicited Comments	0	4	4
Topic Change	0	1	1
Total NCFs	5	8	13
Utterances	9	10	19
DNC	0.56	0.80	0.68

Table 2: Computing the DNC for dialogue (3)

The DNC was computed for all the political interviews in the corpus. Table 3 shows the val-

	Dialogue	Utterances	NCF	DNC
Adversarial Dialogues	1. Paxman v. Howard	54	30	0.56
	Paxman (IR)	24	13	0.54
	Howard (IE)	30	17	0.57
	2. Paxman v. Galloway	48	15	0.31
	Paxman (IR)	19	7	0.37
	Galloway (IE)	29	8	0.28
	4. O'Reilly v. Hartman	36	9	0.25
	O'Reilly (IR)	15	4	0.27
	Hartman (IE)	21	5	0.24
	8. Rather v. Bush (Turns 107-133)	40	19	0.48
Rather (IR)	18	8	0.44	
Bush (IE)	22	11	0.5	
Cooperative Dialogues	7. Keating v. Thatcher	57	0	0
	Keating (IR)	12	0	0
	Thatcher (IE)	45	0	0
	8. Brodie v. Blair (Turns 7-20)	31	2	0.06
	Brodie (IR)	9	1	0.11
	Blair (IE)	22	1	0.05

Table 3: DNC of political interviews in the corpus

ues obtained. Adversarial interviews have a large number of NCFs, thus a high value for the DNC. On the other hand, collaborative exchanges have low occurrence of NCFs (or none at all)⁷.

4 Discussion

There have been previous approaches to modeling dialogue on the basis that participants are not always fully cooperative. Jameson (1989) presents an extensive study for modeling bias, individual goals, projected image and belief ascription in conversation. User-model approaches are flexible to account for intricate situations but, as noted by Taylor et al. (1996), can lead to problems like infinite regress in nested beliefs. Taylor (1994) addressed non-cooperative dialogue behaviour by implementing CYNIC, a dialogue system able to generate and recognise deception; a notion of non-cooperation weaker than the one we address.

More recently, Traum (2008) brought attention to the need for computational accounts of dialogue situations in which a broader notion of cooperation is not assumed: e.g., intelligent tutoring systems, bargaining agents, role-playing training

⁷These results and the validity of DNC measure need further evaluation. We are currently performing two studies: one to determine inter-annotator agreement of the coding scheme for NCFs, and another to test how NCFs correlate to human judgements of non-cooperative conversational behaviour.

agents⁸. Traum's work on conflictive dialogue is mainly aimed at creating virtual humans with abilities to engage in adversarial dialogue. Traum et al. (2008) present a model of conversation strategies for negotiation, that includes variables representing trust, politeness and emotions, and a set of conversational strategies. Despite being adversarial in nature, the conversational scenarios are modeled by means of rules, that are followed by the interlocutors, according to the values of some of the variables. Hence, the dialogues are adversarial, but cooperative under our characterisation of linguistic non-cooperation, and it is not clear how effectively the model accounts for cases in which participants fail to follow the rules of a scenario.

4.1 Working Hypothesis

Finding a suitable model of non-cooperative dialogue involves bridging the gap between the theoretical aspects mentioned so far and the evidence in the empirical data of the previous section. Following Traum and Allen (1994), we base on the hypothesis that non-cooperative features result from decisions that participants make during the conversation, by considering the obligations imposed by the social activity and their individual goals, with an adequate configuration of the priorities for goals and obligations.

Thus, a participant with high priorities for individual goals might compromise the workings of a conversation by choosing contributions that go against the norms of the social activity. On the other hand, participants with higher priorities associated with obligations will favour contributions consistent with the rules of the social activity.

4.2 Research Methodology

For the next steps of the project, we will construct a model based on the hypothesis and test it by means of simulation⁹.

The construction of the model is a formalization of the working hypothesis, including rules for political interviews, goals, obligations, priorities and a dialogue management component. At the

⁸Traum also provides a list of "behaviours of interest", along the lines of the NCFs we identified above: e.g., unilateral topic shifts or topic maintenance, unhelpful criticism, withholding of information, lying, deception, antagonism.

⁹The use of simulation in dialogue modeling was pioneered by Power (1979). It suits our project better than alternatives (e.g., Wizard-of-Oz, dialogue systems), by making it easier to introduce modifications, do re-runs, and generate a large number of cases with different parameter settings.

moment of writing, we are investigating the line of research on obligation-driven dialogue modeling, initiated by Traum and Allen (1994) and developed further by Poesio and Traum (1998) and Kreutel and Matheson (2003).

For the simulation, DPs will be autonomous conversational agents with a cognitive state consisting of goals, a notion of their expected behaviour in a political interview, priorities, and some knowledge of the world. We are currently implementing a prototype based on EDIS (Matheson et al., 2000).

5 Conclusions

In this paper we presented an attempt to shed light on non-cooperation in dialogue by proposing a practical measure of the degree of linguistic non-cooperation in political interviews and a methodology towards a suitable computational model.

Acknowledgments

We would like to thank the NLG group at The Open University (especially Paul Piwek, Richard Power and Sandra Williams) for helpful discussion and comments on previous versions of this paper; and three anonymous reviewers for thoughtful feedback and suggestions.

References

- J.F. Allen and L.K. Schubert. 1991. The TRAINS project. TRAINS Technical Note 91-1. *Computer Science Dept. University of Rochester*.
- S. Attardo. 1997. Locutionary and perlocutionary cooperation: The perlocutionary cooperative principle. *Journal of Pragmatics*, 27(6):753–779.
- N. Blaylock and J. Allen. 2005. A collaborative problem-solving model of dialogue. In *Proceedings of the 6th SIGdial Workshop on Discourse and Dialogue*, pages 200–211, Lisbon, Portugal.
- Harry Bunt. 1994. Context and dialogue control. *THINK Quarterly*, 3.
- H.H. Clark and E.F. Schaefer. 1989. Contributing to discourse. *Cognitive science*, 13(2):259–294.
- P.R. Cohen and H.J. Levesque. 1991. Confirmations and joint action. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, pages 951–957.
- J. Ginzburg. 1996. Interrogatives: Questions, facts and dialogue. *The handbook of contemporary semantic theory*, 5:359–423.
- H. P. Grice. 1975. Logic and conversation. *Syntax and Semantics*, 3:41–58.
- B.J. Grosz and C.L. Sidner. 1990. Plans for discourse. *Intentions in communication*, pages 417–444.
- J. Heritage. 1998. Conversation analysis and institutional talk. Analyzing distinctive turn-taking systems. In *Proceedings of the 6th International Congress of IADA, Tubingen, Niemeyer*.
- A. Jameson. 1989. But what will the listener think? Belief ascription and image maintenance in dialog. *User Models in Dialog Systems. Springer-Verlag*, pages 255–312.
- J. Kreutel and C. Matheson. 2003. Incremental information state updates in an obligation-driven dialogue model. *Logic Journal of IGPL*, 11(4):485.
- C. Matheson, M. Poesio, and D. Traum. 2000. Modelling grounding and discourse obligations using update rules. In *Proceedings of the 1st NAACL conference*, pages 1–8, San Francisco, CA, USA.
- M. Poesio and D. Traum. 1998. Towards an axiomatization of dialogue acts. In *Proceedings of the Twente Workshop on the Formal Semantics and Pragmatics of Dialogues*, pages 207–222.
- R. Power. 1979. The organisation of purposeful dialogues. *Linguistics*, 17:107–152.
- C. Reed and D. Long. 1997. Collaboration, cooperation and dialogue classification. *Working Notes of the IJCAI97 Workshop on Collaboration, Cooperation and Conflict in Dialog Systems, IJCAI 97*, pages 73–78.
- H. Sacks, E.A. Schegloff, and G. Jefferson. 1974. A simplest systematics for the organization of turn-taking for conversation. *Language*, pages 696–735.
- J.R. Searle. 1979. A Taxonomy of Illocutionary Acts. *Expression and meaning: studies in the theory of speech acts*, pages 1–29.
- J. A. Taylor, J. Carletta, and C. Mellish. 1996. Requirements for belief models in cooperative dialogue. *User Modeling and User-Adapted Interaction*, 6(1):23–68.
- J.A. Taylor. 1994. *A multi-agent planner for modelling dialogue*. Ph.D. Thesis, School of Cognitive and Computing Sciences, University of Sussex.
- D.R. Traum and J.F. Allen. 1994. Discourse obligations in dialogue processing. In *Proceedings of the 32nd annual meeting of ACL*, pages 1–8. Morristown, NJ, USA.
- D. Traum, W. Swartout, J. Gratch, and S. Marsella. 2008. A virtual human dialogue model for non-team interaction. *Recent Trends in Discourse and Dialogue. Springer*.
- D. Traum. 2008. Extended Abstract: Computational Models of Non-cooperative dialogue. In *Proceedings of LONDIAL 2008, the 12th Workshop on the Semantics and Pragmatics of Dialogue*, pages 11–14, London, UK.
- D. Walton and E. Krabbe. 1995. *Commitment in dialogue: Basic concepts of interpersonal reasoning*. State University of New York Press.