

Learning to apply Theory of Mind

Rineke Verbrugge (l.c.verbrugge@rug.nl)

Institute of Artificial Intelligence, University of Groningen, Grote Kruisstraat 2/1. 9712 TS Groningen, The Netherlands

Lisette Mol (l.mol@uvt.nl)

Department of Humanities, Communication and Cognition, Tilburg University, Warandelaan 2, 5037 AB, Tilburg, The Netherlands

January 23, 2008

Abstract. In everyday life it is often important to have a mental model of the knowledge, beliefs, desires, and intentions of other people. Sometimes it is even useful to have a correct model of their model of our own mental states: a second-order Theory of Mind. In order to investigate to what extent adults use and acquire complex skills and strategies in the domains of Theory of Mind and the related skill of natural language use, we conducted an experiment. It was based on a strategic game of imperfect information, in which it was beneficial for participants to have a good mental model of their opponent, and more specifically, to use second-order Theory of Mind. It was also beneficial for them to be aware of pragmatic inferences and of the possibility to choose between logical and pragmatic language use. We found that most participants did not seem to acquire these complex skills during the experiment when being exposed to the game for a number of different trials. Nevertheless, some participants did make use of advanced cognitive skills such as second-order Theory of Mind and appropriate choices between logical and pragmatic language use from the beginning. Thus, the results differ markedly from previous research.

Keywords: Theory of Mind, Epistemic Logic, Strategic Reasoning, Pragmatics, Skill Acquisition

1. Introduction

In everyday life, people frequently make use of their ability to reason about others and to infer the implicit meaning of sentences. Consider the following two situations:

Situation 1 You are called by a friend who asks you for a phone number. You know the number by heart, so you ask her whether she has pen and paper. She answers you with “No, I don’t”. Can you conclude that she also does not have a pencil and paper ready?

Situation 2 You are playing happy families¹ and you are the first to pose a question. You ask your opponent for the ‘elephant’ of the family ‘mammals’. Your opponent replies with “No, I don’t have this card”. Can you conclude that he doesn’t have any member of the mammals family?

¹ In some countries this game is called ‘quartets’.

In the first case, you *know* that your friend *has the desire* to be cooperative and thus your reasoning would be something like: She does not have a pencil, for if she did she would have told me so, since she knows it is relevant. In the second case you *know* that your opponent *does not want* you to *know* which cards he has, since he has the desire to win the game. You therefore are aware that he would not tell you whether he has any other members of the family, unless he really had to, so you do not conclude that he does not have them.

These examples point out that people use their knowledge about the situation and about others to determine the meaning of a sentence. It would be interesting to know how humans use and acquire such skills. In the study described in this article, we investigated to what extent adults use and acquire complex skills in the domains of reasoning about others and language use. We are particularly interested in *Theory of Mind*, the cognitive capacity to understand and predict external behavior of others and oneself by attributing internal mental states, such as knowledge, beliefs, and intentions (Premack and Woodruff, 1978). Readers of this journal, being well-trained in logic, are probably quite capable of using nested forms of Theory of Mind and have little difficulty interpreting sentences like “Johan does not know that Alice knows that it is not common knowledge among logicians that Patrick has written a novel under pseudonym” or analyzing games that require sophisticated forms of recursive reasoning about strategies. Introspection suffices to know that, at least in some circumstances, some people can reason correctly at various levels of Theory of Mind, and no experiment can deny this. However, for a more general question such as “under what conditions do people generally engage in Theory of Mind, to which degree do they apply it correctly, and can they learn to apply it in an unusual context?”, empirical research is needed. This article presents some experimental findings indicating that the degree to which people correctly apply Theory of Mind, may be rather less than is often assumed.

Many empirical studies of human reasoning involving formal models like logic have brazenly concluded that logics are not at all adequate as models for human reasoning (see for example (Wason, 1966; Johnson-Laird, 1983)). Our stance towards these formal models is much friendlier. We find that logics help at least to formulate empirical questions precisely, for example by clearly delineating different orders of Theory of Mind (e.g. in (Parikh, 2003)), and we gratefully use logics in this modest way in the current article. It is our ambitious aim for future research to create flexible formal logical models that are tailored, more closely than the currently available epistemic logics, to individual human capabilities of applying Theory of Mind.

As a focus of our experiments, we use the highly specialized context of a competitive new game of imperfect information called Mastersminds. In this adversarial context, in contrast to the cooperative context usually used in discourse understanding experiments, learning seems bound to be necessary.

More specifically, both learning to strategically apply second-order Theory of Mind and learning to produce and interpret uncooperative communication are expected.

In the next section, we describe a number of theories relevant to the present study. Section 3 describes our research question and hypotheses. To test these hypotheses, we conducted an experiment in which participants played a game against each other. This game was a variant of the game *Mastermind*. We describe the experimental setup and the predictions, followed by the results of the experiment. Then follow a discussion of the results and the conclusions that can be drawn from this study. In the last section, we present some ideas for future work.

2. Background: Formal Models for Real People

Several research enterprises have been combined in this study. First of all, we will give a reminder of the empirical investigation of Theory of Mind and its formal counterpart, the logic of knowledge (epistemic logic). Then we shortly describe game theory and its application to the game of *Mastersminds* used in our experiments. This is followed by a description of pragmatic and logical interpretations of scalar terms, which will play a role in the feedback sentences employed in *Mastersminds*. We close the section with a short overview of different theories about the interplay between controlled and automated processes in the acquisition of all complex skills relevant for our experiments: higher-order Theory of Mind, strategic game playing, and pragmatic and logical language use.

2.1. THEORY OF MIND USE

One of the advanced skills that we are interested in is the use of Theory of Mind (ToM). Although normally developing children from the age of around five are able to distinguish between their own mental states and those of others, Keysar et al. (Keysar et al., 2003) argue that even adults do not reliably use this sophisticated ability to interpret the actions of others. They found a stark dissociation between the ability to reflectively distinguish one's own beliefs from others', and the routine deployment of this ability in interpreting the actions of others. The second didn't take place in their experiment. In other experiments by the same research group, similar results were found (Keysar et al., 2000; Keysar et al., 1998; Horton and Keysar, 1996).

To have a first-order ToM is to assume that someone's beliefs, thoughts and desires influence one's behavior. A first-order thought could be: 'He does not know that his book is on the table'. In second-order ToM it is also recognized that to predict others' behavior, the desires and beliefs that they

have of one's self and the predictions of oneself by others must be taken into account. So, for example, you can realize that what someone expects you to do will affect his behavior. For example, '(I know) he does not *know* that I *know* his book is on the table' would be part of my second-order ToM. To have a third-order ToM is to assume others to have a second-order ToM, etc.

In defining the different orders, two choices have been made. The first is that to increase the order, another agent must be involved. 'I know his book is on the table' and 'I know I know his book is on the table' are said to be of the same order. A motivation for this choice is that these statements are equivalent in the system S5, which is used in epistemic logic (see subsection 2.2). So for the order to increase, the agents that the knowledge is about must be different. Here we abstract away from the discussion in (Nisbett and Wilson, 1977; Dunin-Kępicz and Verbrugge, 2006) that in reality, people do not perfectly reason about their own knowledge and ignorance. Thus epistemic logic is an idealized model with respect to human powers of reflective introspection, too, but in this article we concentrate on people reasoning about one another.

The second choice we made is to consider both 'I know p ' and p to be zeroth-order knowledge for atomic propositions p . This mainly is a matter of speech. The fact p in itself, which can be true or false, only becomes knowledge when it is known by someone. So only when someone knows that p , p can be considered zeroth-order knowledge.

2.2. EPISTEMIC LOGIC

Epistemic logic is a type of modal logic that can be used to describe the knowledge and beliefs of an agent, or a system of agents. In epistemic logic, the K_i operator represents that agent i knows something. For example, K_1p means *agent 1 knows p* . By definition, an agent can only know things that are true, reflected in the 'truth axiom' $K_i\phi \rightarrow \phi$ for all formulas ϕ .

The K_i operator can take scope over an epistemic formula. For example, $K_1(p \rightarrow q)$ for *agent 1 knows that p implies q* , or K_1K_2p for *agent 1 knows that agent 2 knows that p* . Especially the last example is of interest here. By nesting of the modal operator K_i , knowledge from different orders of Theory of Mind can be represented.

This is relevant to describe knowledge of agents playing Mastermind, a game of which we use a variant. Mastermind is a game in which player 2 has to guess the secret code of four colors, that is composed by player 1. In reply to each guess made by player 2, player 1 has to specify how many colors from that guess match colors in the secret code, and how many of these are in the right place.

The fact that agent 1 has the first-order knowledge that agent 2 knows that red occurs in agent 1's secret code of four colors could be represented by K_1K_2p , where p means *Red occurs in the secret code of agent 1*. In a similar

way, $K_1K_2K_1p$ would mean that agent 1 knows that agent 2 knows that agent 1 knows that red is in his secret code. This knowledge that agent 2 knows that agent 1 knows that p is part of the second-order Theory of Mind of agent 1. The first K_i operator only specifies whose knowledge is represented, whereas the ones following are part of that knowledge itself.

The standard system S5 of epistemic logic, in addition to the truth axiom mentioned above, contains two more axioms. $K_i\phi \rightarrow K_iK_i\phi$ ('positive introspection') expresses that agents are aware of what they know. On the other hand, $\neg K_i\phi \rightarrow K_i\neg K_i\phi$ ('negative introspection') expresses that agents are aware of their ignorance. As mentioned before, these reflective properties make for a rather idealized formal model of human reasoning.

In addition to the K_i operator, the dual M_i operator represents what an agent thinks that might be (i.e. what is consistent with his information), and is defined by $M_i\phi \Leftrightarrow \neg K_i\neg\phi$. The B_i operator represents what an agent believes, the D_i operator what an agent desires, and the I_i operator what an agent intends. When considering a finite system of multiple agents, two more operators are of interest. E , for *every one knows that*, and C for *it is common knowledge that*. Agents are said to have common knowledge of p if it is the case that everyone knows that p (so Ep), everyone knows that Ep , EEp , etcetera ad infinitum. For more on epistemic logic and its possible worlds semantics, see (Fagin et al., 1995; Meyer and van der Hoek, 2004; Van der Hoek and Verbrugge, 2002).

2.3. STRATEGIC GAMES

According to game theory, a game is being played whenever people interact with each other. Game theory provides a way to formally describe and categorize games and strategies. It has successfully been applied to problems in economics, political science, biology and social philosophy. A *strategy* in game theory is defined in a formal way. To understand strategies, let us first look at games. A game must have rules, which specify for each player what can be done at what point in time. In addition, the rules should specify how the players are rewarded at the end of the game.

A game can be represented as a game tree. The root of the tree represents the first move of the game; for example, in figure 1, the root is the leftmost node at the top. The leaves correspond to possible game endings and should be labeled with the outcomes of the game for each player; for example in figure 1, the rightmost leaf at the bottom is labeled 2,4 corresponding to an outcome of 2 for player 1 and 4 for player 2. Each node represents a possible move and the player who is to move, while the edges leading away from a node represent the actions possible at that point in the game; for example, the root in figure 1 is labeled 1 to represent that it is player 1's turn and there are two edges leading out, which are labeled with capital letters corresponding to

player 1's possible choices, either to go R (to the Right) or D (Down); player 2's choices at the second node are labeled with the corresponding lowercase letters. A play then consists of a connected chain of edges, from the root to one of the leaves. A *pure* strategy for player i can now be defined as 'a statement that specifies an action at *each* of the decision nodes at which it would be player i 's duty to make a decision if that node were actually reached' (Horton and Keysar, 1996).

The knowledge that is relevant to the game, that a player has at a particular point in time is represented in game theory as an *information set*. Information can be either *perfect* or *imperfect*. In bridge, information is imperfect, because players do not know the hands of other players. This is different from the game of chess, a perfect information game where each player always knows everything there is to know about the current state of the game, because both players can see the board and the positions of the pieces. In a game of imperfect information, a pure strategy needs to specify what action a player performs at each possible information set (see (Binmore, 1992) for more on game theory).

The game used our experiment, a symmetric version of Mastermind, is far more complex. It has been designed by Grachten and Druiven and described by (Kooi, 2000) and is called Mastersminds. We have used a variant that makes it suitable for our experiment. Mastersminds is a two-player *zero-sum* game. *Zero-sum* means that the pay-offs of all players always sum to zero: if one player wins, the other player loses, so the game is strictly competitive. In our variant of Mastersminds, each player has a secret code of four different, ordered colors, and has to guess the secret code of the opponent in order to win. This can be done by making guesses and receiving feedback sentences on how many colors are right, and how many of them are in the right position. Each turn, one player makes a guess about the other player's array of colors. In addition to receiving feedback sentences, this player also evaluates *this same guess* about his own array of colors and sends this information to his opponent by selecting sentences.

The game tree for Mastersminds would become extremely large, because there are many possible ways for the game to evolve. In addition, the information players have depends on what choices their opponent makes, which results in many possible information sets. Therefore, it would be very impractical to define a formal strategy for this version of Mastersminds: already for the classical asymmetric Mastermind, nobody knows a winning strategy (but see (Kooi, 2005)). However, global strategies for Mastermind can be described, such as 'revealing little information' or 'concentrating on guessing the opponent's secret code'. In the empirical part of the article, *strategy* is used in this less formal way.

2.4. GAME THEORY AND REAL PEOPLE

In classical game theory, it is often assumed that there is common knowledge that players are rational and capable of perfect reasoning. For Mastersminds, this is not a plausible assumption. It would imply that all players would be capable of unboundedly high order ToM reasoning. Game theory has often been criticized for the assumption of perfect reasoning. In fact, many experimental studies have shown that people do not always follow rational strategies, for example, they do not do so in so-called centipede games (introduced in (Rosenthal, 1981)). Let us remind the reader of these games of perfect information, using the small example of figure 1.




Figure 1. Game tree for the centipede game

In this centipede game, two players take turns choosing either to take a larger share of the current amount of marbles, namely two more than the opponent (down in the picture), thereby ending the game; or alternatively, to pass the choice to the other player (right in the picture), which leads to an increase of the total available amount of marbles by two. In the picture, The pay-offs for each player are represented at the leaves, and the amount of marbles at the start of the game is one. If player 1 starts choosing ‘down’, he receives one marble while player 2 receives nothing; if, on the other hand, player 1 chooses ‘right’, the available amount increases by two. At the next turn, if player 2 chooses ‘down’, she receives two marbles while player 1 receives nothing, and so on.

Now we can demonstrate the relevance of common knowledge of rationality by using backward induction. Let r_1 denote that player 1 is rational, and r_2 that player 2 is rational. At the fourth and last choice point from the left, using r_2 , we can infer that 2 will prefer four marbles to three ones and she will choose ‘down’. Since rationality of both players is common knowledge, we know that $K_1 r_2$, hence, when in the third choice point, player 1 knows that player 2 will choose ‘down’ in the fourth choice point, and hence, since player 1 is rational and prefers three marbles over two ones, he will choose ‘Down’. Since $K_2 K_1 r_2 \wedge K_2 r_1$, if player 2 were to reach the second choice point from the left, she would apply the same reasoning that we just did and

conclude that player 1 will play ‘Down’ in the third choice point, so player 2, being rational, will play ‘down’ in the second choice point. Continuing this line of reasoning, and using the fact that $K_1K_2K_1r_2 \wedge K_1K_2r_1$, we can conclude that player 1 will play ‘Down’ at the start. Therefore, if there were common knowledge of rationality, then backward induction, which is based on high orders of Theory of Mind, could be used by the first player in a centipede game to conclude that he should immediately opt for the first dead-end and stop the game (see also (Van der Hoek and Verbrugge, 2002)).

It seems paradoxical that reasoning on the basis of common knowledge of rationality leads to a less than optimal outcome for both players. Indeed, empirical research has shown that instead of immediately taking the ‘down’ option, players often show partial cooperation, moving right for several moves before eventually choosing to take the down option (McKelvey and Palfrey, 1992; Nagel and Tang, 1998). Nagel and Tang suggest as possible reason for this deviation from the game-theoretic outcome that players sometimes have reason to believe that their opponent could be an altruist who always cooperates (by moving to the right). In such a case, it is better to join the altruist in going to the right and then defect on the last round. In (Van Eijck and Verbrugge, 2008), Van Benthem and Van Eijck present a Platonic dialogue on *Game theory, logic, and rational choice* that offers an intriguing and sophisticated logical analysis of backward induction and its rival strategies.

Another possible explanation of human behavior in the centipede game involves error and cognitive limits: if the opponent has not correctly performed the full backward induction, it may be advantageous to cooperate in the first rounds. On this line of cognitive limits, Hedden and Zhang, Colman and many others have shown that in reality, players hardly use backward induction at all and make use of only first-order or second-order ToM (Hedden and Zhang, 2002; Colman, 2003). Hedden and Zhang used a game of perfect information very similar to the centipede game and found that adult subjects would start using at most first-order theory of mind. Gradually, a number of subjects would shift to second-order theory of mind when they started modeling their opponent as a first-order reasoner (Hedden and Zhang, 2002) (but see (Flobbe et al., 2007) in this issue for an alternative interpretation and experiments).

It would be interesting to see if similar learning to apply higher orders of ToM as proposed by Hedden and Zhang occurs in Mastersminds, our game of imperfect information. It is known that backward induction cannot be applied to games of imperfect information, so the formal model of computing an optimal solution consisting of pure strategies (technically, a subgame-perfect equilibrium) by backward induction as described above, does not apply to Mastersminds (Binmore, 1992).

2.5. PRAGMATIC INFERENCES

Besides ToM reasoning and strategic reasoning, a third skill that we investigate is language use, especially drawing pragmatic inferences. According to Grice (Grice, 1989), people use the quantity maxim to infer the implicit meaning of a sentence. The quantity maxim states that interlocutors should be as informative as is required, yet not more informative than is necessary.

Using the quantity maxim it can be inferred that, for example, if a teacher says ‘Some students passed the test’, it is the case that not all students passed the test. This is because if all students would have passed the test, the teacher would have used the more informative term *all* instead of the weaker term *some*, since otherwise the quantity maxim would have been violated.

Some and *all* are scalar terms. Scalar terms can be ordered on a scale of pragmatic strength. A term is said to be stronger if more possibilities are excluded. An example is $\langle a, \textit{some}, \textit{most}, \textit{all} \rangle$ which is ordered from weak to strong. The above phenomenon of inferring ‘not all students’ from the use of ‘some students’ is an example of a scalar implicature. In case of a scalar implicature, it is communicated by a weaker claim (using a scalar term) that a stronger claim (using a more informative term from the same scale) does not hold.

Feeney et al. (Feeney et al., 2004), propose that there are three stages to people’s understanding of *some*:

- (a) the logical (truth-conditional) interpretation which precedes children’s sensitivity to scalar implicatures;
- (b) the pragmatic interpretation which results from drawing pragmatic inferences;
- (c) a logical interpretation that results from choice rather than from the incapability to make the pragmatic inference.

The first two stages are in line with the results in (Noveck, 2001) and (Papafragou and Musolino, 2003). Feeney et al. found evidence for a third stage, in which adults can choose a logical interpretation over a pragmatic interpretation, even though they can make the pragmatic inference that *some* implies *not all*. They conducted an experiment in which undergraduate students performed a computerized sentence verification task. They recorded the students’ answers and reaction times. Here are two of the *some* sentences they used.

1. Some fish can swim.
2. Some cars are red.

In contrast to the second sentence, the first sentence is *infelicitous*: world knowledge contradicts the pragmatic conclusion that not all fish can swim, so that the logical and pragmatic interpretation are in conflict. Feeney et al. found that for participants who gave logical responses only, reaction times for responses to infelicitous *some* sentences such as 1 were longer than those for logically consistent responses to felicitous *some* sentences as 2. Notice that to both sentences the logical response is ‘true’. The pragmatic response to 2 is ‘true’ as well. The pragmatic response to 1 is ‘false’. So the sentences in which the logical and pragmatic response are in conflict resulted in longer reaction times. These results favor a theory that logical responses are due to inhibition of a response based on the pragmatic interpretation over a theory that logical responses result from failure to make the pragmatic inference. This suggests that a more logical language use in adults can be seen as an advanced skill.

An interesting explanation of scalar implicatures based on bidirectional optimality theory has been presented in (Blutner, 2000). Moreover, Dekker and Van Rooij have convincingly argued for a correspondence between pragmatic inferences according to bidirectional optimality theory and strategic reasoning in games (Dekker and Van Rooij, 2000) (see also (Flobbe et al., 2007; Mol, 2004)). Finally, Benz and Van Rooij convincingly show that in a cooperative dialogue, optimal assertions are calculated by using backward induction (see subsection 2.3), accounting in a uniform way for scalar implicatures as well as for several other types of implicature (Benz and van Rooij, 2007).

The kinds of discourse usually considered in formal semantics—and the above-mentioned articles are no exception—concern cooperative communication. We feel that it is at least as important to investigate adversarial communication, where even if one has the required information, one provides the interlocutor with as little information as one can get away with. Our analysis of the subjects’ discourse in the experimental study focuses on the contrast between cooperative and adversarial communication. In our model, cooperative communication corresponds with Gricean pragmatics, whereas adversarial communication corresponds more closely to classical logic ².

² Note that there are alternative viable approaches to this contrast, for example using defeasible, non-monotonic logics. There, cooperative communication leads to ‘credulous’ inferences (where the conclusion holds in at least one extension, an intended model of the premises) whereas adversarial communication leads to ‘skeptical’ inferences (where the conclusion should hold in all possible extensions) (Stenning and van Lambalgen, 2007; Antoniou, 1997).

2.6. LEARNING BY REFLECTION

Human reasoning has often been divided into automatic and controlled processing. Automatic reasoning is unconscious and is usually modeled in a connectionist architecture. Controlled reasoning is conscious and is usually modeled in symbolic systems. In (Schneider and Chein, 2003; Schneider and Shiffrin, 1977), Schneider et al. describe automatic processing as the activation of a sequence of nodes that “nearly always becomes active in response to a particular input configuration, [...] without the necessity for active control or attention by the subject”. Automatic processes “require an appreciable amount of consistent training to develop fully”.

Controlled processes on the other hand are described as “a temporary sequence of nodes activated under control of, and through attention by, the subject.” They are “tightly capacity limited, but the costs of this capacity limitation are balanced by the benefits deriving from the ease with which such processes may be set up, altered, and applied in novel situations for which automatic sequences have never been learned” (Schneider and Shiffrin, 1977).

Classical theory of skill acquisition (Fitts, 1964) describes learning as a process of automation: one starts a new skill in the cognitive stage, in which controlled deliberate reasoning is needed to perform the task. This stage is characterized by slow performance and errors. By repeatedly performing the skill, eventually the autonomous stage is reached, where performance is fast and automatic, requiring little working memory capacity. During the process of automation, the control that one has over the process of performing the task decreases. Deliberate access to automated skills is therefore limited.

Although this theory can explain many phenomena, it is limited: Skills are usually considered in isolation, whereas in reality they build on one another. For example, the skill of multiplication is based on the skill of addition. However, according to Fitts’ theory, mastered and hence automated skills cannot in themselves serve as a basis for more advanced skills, because deliberate access to automated skills is limited. Hence, in this classical theory it remains unclear how transfer of knowledge from one skill to another is possible.

In (Karmiloff-Smith, 1992), however, it is reported that children can only describe what they are doing after they have mastered a skill (e.g., in number conservation experiments). Thus, the capacity for deliberate reasoning sometimes increases rather than decreases with expertise. This cannot be explained by assuming skill acquisition to end in the autonomous stage. In (Sun and Zhang, 2004), Sun and Zhang stress the interaction of implicit (automatic) and explicit (controlled) reasoning processes during skill acquisition. In (Evans, 2003), Evans argues that most reasoning tasks have automatic and deliberate components. Other than suggested by Fitt’s theory, the two types of reasoning thus seem to be closely intertwined.

We suggest that skill acquisition is a continuous interplay between deliberate and automatic processes. During the initial stages of skill acquisition all deliberation is focused on basic performance of the task. Once performing the task becomes more automatic, deliberate processes can shift to reflection on the task. Basic performance of the task and reflection on the task can be considered dual tasks, both competing for resources. Once performance becomes more automated, more resources become available for reflection. Reflection allows using the skill as a building block for more complex skills, looking ahead a few steps, or, as is the case in the present study, reasoning about others' knowledge (also see (Taatgen, 2005) for examples of how automation can improve flexibility in reasoning).

We assume that to reach expert level performance in domains such as reasoning about others, pragmatics, and learning from instruction, deliberate reasoning processes, such as self-monitoring, are crucial.

3. Research Question and Hypotheses

The context described in the previous section leads to the following problem statement: *How do deliberate and automatic processes interact in the acquisition of complex skills?* This study is an exploratory investigation, for which we have stated the following research question: *To what extent do people use and acquire complex skills and strategies, in the domains of reasoning about others and language use?* We have narrowed this down to the specific case of playing Mastersminds, the symmetric version of the game Mastermind, introduced in subsection 2.3. We make use of three hypotheses:

HYPOTHESIS 1. *Performing a task and simultaneously reflecting upon this task can be seen as a form of dual tasking.*

This hypothesis states that when people perform a task that involves reasoning with incomplete information, or drawing pragmatic inferences, reflection can be considered a secondary task. The first task includes reasoning based on one's own knowledge and the truth-conditional (e.g. logical) meaning of utterances. The second task is more complex, and includes using reflection to reason about others and to infer from pragmatically implicated meaning. These tasks compete for resources and their demands decrease with skill acquisition.

When playing Mastersminds (see subsection 2.3 and section 4), the first task is playing the game according to its rules. This involves reasoning about the game rules and determining which sentences are true. The second task is developing a winning strategy. This involves reasoning about what the opponent thinks, is trying to make you think, or thinks that you are trying to

make him think, etc., as well as determining what is pragmatically implicated by an utterance, or which utterances are not only true but also reveal the least information.

HYPOTHESIS 2. *In an uncooperative conversation, people will shift their interpretation and production of quantifiers from a pragmatic use (according to Grice's quantity maxim) to a less pragmatic use (violating Grice's quantity maxim).*

The idea here is that in an uncooperative situation, people will be aware that others are trying to reveal little information (part of first-order ToM) and therefore will not be as informative as possible (as the quantity maxim requires). They will therefore not use the pragmatic inferences that they usually do in interpretation. In addition, people may develop more logical productions to be less informative themselves.

The reasoning necessary for this change in strategy is part of the secondary task of reflective cognition. Therefore, people will only be able to make this change when the first task is sufficiently automated.

HYPOTHESIS 3. *In using quantifiers, people make use of an automated process, which results in a pragmatic use of the quantifier. This automated process can be 'overruled' by a deliberate reasoning process, which results in a logical use of the quantifier.*

Hypothesis 3 is on what kind of reasoning is involved in using quantifiers, especially to make the shift described in hypothesis 2. The theory of three stages that is proposed in (Feeney et al., 2004) (see section 2.5) seems in line with the theory of skill acquisition we propose. If so, the process of making pragmatic inferences should be an automated process in adults and the ability to overrule this pragmatic interpretation would result from reflective cognition. Once the demands of the first task have decreased sufficiently, this reflective cognition can take place, resulting in the change of strategy described in hypothesis 2.

4. Master(s)Mind(s) Experiment

4.1. PARTICIPANTS

All twelve participants were at least 18 years old and had not followed any courses on the formal theories relevant for the experiment such as epistemic logic, game theory or linguistic theories. They were native speakers of Dutch.

4.2. DESIGN

Participants had to complete two sessions, each of about three hours, in which they played a symmetric head to head game via connected computers in the same room. Participants were facing a different opponent in each session. In this game they had to correctly guess the secret code, consisting of four different, ordered colors, of their opponent. Players gave each other feedback by selecting Dutch sentences from a list. Although not explicitly told to participants, these sentences differed in pragmatic strength. As proficient players of the game would discover, the game was about gaining as much information as possible, while at the same time revealing as little information as possible. Because of this second aspect, the conversation is not fully cooperative and thus hypothesis 2 is relevant.

During the game, players had to submit their interpretation of the sentences they received as feedback. They had to submit all the ‘worlds’ that they thought to be possible given the feedback sentences, using a code: For each right color in the right position they had to select a black circle and for each color which was correct but in the wrong place, a white circle. To represent ambiguity and vagueness, participants could submit more than one combination of black and white circles that they considered possible.

Let us look at an example. Imagine John having the secret code 1 = red, 2 = blue, 3 = green, 4 = yellow and Mary guessing 1 = red, 2 = orange, 3 = yellow, 4 = brown. The evaluation of this situation is that exactly one guessed color (namely red) is right and in the right place and exactly one guessed color (namely yellow) is right, but in the wrong place. John has to choose two feedback sentences to send to Mary, one about color and one about position. He could communicate:

‘Sommige kleuren zijn goed.’ (‘Some colors are right.’)

and

‘Een kleur staat op de goede plaats.’ (‘There is a color which is in the right place.’)

This would indicate that John thinks that *sommige* (some) can be used to refer to *exactly two* and that *een* (a) can be used to refer to *exactly one*. This is a pragmatic production: in accordance with Grice’s maxims. If John had instead chosen the sentence ‘Een kleur is goed’ (‘A color is right’), then he would allow *een* (a) to refer to a situation where *exactly two* holds. This would be a more logical production (in logic *a* is true in case of *at least one*). Notice that because the sentences were communicated in written form, through the computer program, intonation and non-verbal communication were absent.

Mary now has to give her interpretation of the sentences chosen by John. So if she thinks that, given the first two sentences, it could be the case that two colors are right, of which one is in the right position, she would submit

(black, white) as a possible interpretation. If she considers the situation where three colors are right, of which two colors are in the right position, possible as well, she would also submit (black, black, white). If she would only submit the first possibility, her interpretation would be pragmatic. If she would also submit the second one, her interpretation would be more logical.

In the experiment Mary would have to give John feedback about her guess compared to her own secret code as well, and John would then submit his interpretation of those sentences. Each turn, one player can make a guess, in this example Mary.

During the experiment participants had to answer questions. The purpose of those questions was to get information on their strategy and the order of the Theory of Mind they were using. For the same purpose, participants completed a questionnaire after each session. More details on the experiment can be found in (Mol, 2004).

4.3. PREDICTIONS

Since players of the game *Mastersminds* have to perform many actions each turn, we expect participants to start with a very simple or no strategy. As they get more experienced in playing the game they will have enough resources left for developing a more complex strategy.

Grice's maxims are applicable in situations where conversation is cooperative. Since a rational strategy for playing the game in the experiment is being as uninformative as possible, communication will probably not be cooperative in the experimental conditions. So once the participants have mastered the game well enough to think about strategy and have become familiar with the uncooperative context, we expect them to develop a less pragmatic use of the sentences.

We predict that while playing the game, the order of the theory of mind used by the participants increases. This will lead to the participant considering the amount of information that is revealed by the feedback sentences chosen, and the amount of information that will have to be revealed as a result of a guess made (first-order ToM). The participant will also become aware that his opponent is trying to reveal little information (second-order ToM). This will lead to a more logical interpretation. Eventually, the participant may use the knowledge that his opponent knows that he is trying to hide certain information (third-order ToM).

We also predict that there will be individual differences in what order of ToM will be used and how logical the language use becomes, as well as in the speed of developing a better strategy, because there will be individual differences in available processing capacity. If participants eventually reach logical language use, then this results from a conscious reasoning process.

Therefore, we expect participants to be able to describe this part of their strategy.

4.4. RESULTS

Participants are numbered from 1 to 12. Participants 10, 11 and 12 completed only one three hour session.

Table I. *Highest Order of ToM used.* This table shows the highest order of ToM that participants used during the experiment. The numbers represent the participants. The order used was determined from the answers participants gave to questions that were asked during the experiment.

1st-order	possibly 2nd-order	2nd-order
3, 5, 6, 7, 12 8, 9, 10	4	1, 2, 11

Three out of twelve participants showed clear signs of the use of second-order ToM (table I). One additional participant probably used second-order ToM as well, but in this case it was less clear whether the knowledge was applied. An example of second-order ToM use in this game is that agent 1 makes his guesses based on the assumption that ‘the guesses made by agent 2 are evasive about agent 2’s own code, since agent 2 does not want agent 1 to know agent 2’s secret code’. All of these four participants played in accordance with a strategy to be uninformative (table II) and had a fairly to strict logical language use (table III).

The remaining eight participants all used first-order ToM. An example of first-order ToM use in this game is that agent 1 takes into account what agent 2 already knows about agent 1’s secret code. Two of these participants used the strategy of being uninformative and had a fairly logical language use, similar to the participants who used second-order ToM. The other six used the strategy of being informative or a strategy which did not consider the amount of information revealed and had a fairly to strict pragmatic language use.

Participants that stated to have the strategy of being uninformative and had a fairly to strict logical language use showed a type of behavior which the others did not show (table IV). This behavior consists of preferring less informative sentences to more informative ones in production. For example, favoring sentence 1 over sentence 2 in a case where, from a logical perspective, they both hold.

1. ‘Some colors are right.’

Table II. *Strategy*. This table shows what kind of strategy participants used during the experiment, initially and finally. The strategy was determined from answers that participants gave to questions posed during the experiment. The numbers of the participants who made a shift are in italic in the row that represents the final strategy.

	being uninformative	being informative	other
initially	1, 2, 4, 5, 10, 11	3, 8, 9, 12	6, 7
finally	1, 2, 3, 4, 5, 11	9, 12	2, 3, 6, 7, 8 10

Table III. *Language use*. This table shows the type of productions and interpretations (logical or pragmatic) that participants submitted during the experiment, initially and finally. The numbers represent the participants. The numbers of the participants who made a shift are in italic in the row that represents the final language use.

	pragmatic	fairly pragmatic	fairly logical	logical
initially	8	5, 6, 7, 9, 10, 12	1, 2, 3, 4	11
finally	6, 7, 8, 12	9, 10	1, 2, 3, 4, 5	11

2. 'All colors are right.'

All participants who used second-order ToM did so from the start. No shifts in order of ToM used were observed. Some shifts were measured in language use. One participant shifted from a fairly pragmatic to a fairly logical use. This participant used the strategy to be uninformative. Three participants shifted from a fairly pragmatic to a fully pragmatic use. They did not use the strategy of being uninformative. The other participants were constant in their language use.

One participant shifted from using the strategy of being informative to using the strategy of being uninformative. This participant had a fairly logical language use. One participant abandoned the strategy of being uninformative,

Table IV. *The preference for uninformative sentences.* This table indicates which participants preferred using less informative sentences in production. The numbers represent the participants. The numbers of the participants who made a shift are in italic in the row that represents the final behavior.

	preferred less informative sentences	did not prefer less informative sentences
initially	1, 3, 4, 5, 11	2, 6, 7, 8, 9, 10, 12
finally	1, 2, 3, 4, 5, 11	6, 7, 8, 9, 10, 12

in order to give the opponent a better chance of winning (!). This participant had a fairly pragmatic use of language.

The participants using more advanced strategies clearly had to put in little effort to play the game and understand the computer program used. For example, this could be concluded from the fact that they made very few mistakes in playing the game by the rules, as well as from their answers in the questionnaire about their familiarity with computers and with strategic games. The people with the least advanced strategies made more mistakes in playing the game than others did.

Most participants wrote down thoughts on the meaning of scalar terms, the terms they considered possible and their strategy in their answers to the questions posed during the experiment.

5. Discussion and Conclusions

5.1. BACK TO THE FORMAL MODELS FOR REAL PEOPLE

In interpreting the experimental results, the language of the extended epistemic logic introduced in subsection 2.2 has been beneficial. For example, during the experiment, participant 1 wrote in an answer to a question that in making a guess, he considered that the guesses made by his opponent were evasive from his opponent's secret code, since he wanted to hide this code. Participant 1 was applying the knowledge that his opponent wanted to hide his secret code, in other words, that his opponent did not want him to know his secret code. Let c denote *the secret code of participant 1's opponent is* d , and let participant 1 and his opponent be denoted by agent 1 and agent

2 respectively. Now the fact that participant 1 has this knowledge could be represented as follows in epistemic logic: $K_1 D_2 \neg K_1 c$ ³.

As explained in subsections 2.1 and 2.2, the order of this knowledge can be found by counting the operators, as long as the agents are different and the first operator, which serves to indicate which agent has the knowledge, is left out. The representation of the knowledge participant 1 has (as opposed to the representation of the fact that participant 1 has this knowledge) is: $D_2 \neg K_1 c$. Now there are two operators with different agents, of which the first is not the agent that has the knowledge and thus it can be concluded that this knowledge is part of a second-order ToM.

5.2. COMPLEX SKILLS AND STRATEGIES

The research question for this study was: *To what extent do adults use and acquire complex skills and strategies, in the domains of reasoning about others and language use?* It was found that four out of twelve participants used all complex skills that were measured: a strategy of being uninformative, using logical interpretation and production, and using second-order ToM reasoning. They did so from the start of the experiment. By the end of the experiment, half of the participants had a logical interpretation and production and also half of the participants had the strategy of being uninformative. So the complex skills that we investigated were used to a fair extent.

One participant developed the strategy of being uninformative during the experiment. Another participant developed a more logical language use. A third participant developed a better strategy in production. Altogether, skill acquisition during the course of the experiment was very limited.

There clearly were individual differences. Some participants did not seem to use any complex skills or strategies. Some participants even developed a more pragmatic language use. This may be because the pragmatic meaning of some of the scalar terms that were used depended on the situation, and hence their interpretation wasn't automated yet.

Dual task model

Hypothesis 1 stated that performing a task and simultaneously reflecting upon this task can be considered dual-tasking. Although participants showed little development during the experiment, the results are in line with this hypothesis. Three groups of participants can be discriminated.

The first group consists of the four participants (1, 2, 4, and 11) that showed the use of all complex skills from the start of the experiment. It could be that the first task, playing the game according to its rules, was relatively

³ Here we abuse the language a bit, because in fact participant 1 cannot refer to the real code of participant 2, so the propositional atom c remains vague. For solutions to such problems, see the literature on epistemic logic for security protocols, e.g. (Teepe, 2006).

easy for these participants. They made few or no mistakes, and had relatively much experience in working with computers and playing strategic games, which indeed points in this direction. Because the cognitive demands of the first task were small for this group, they could use a lot of their resources for the second task: reasoning about others. Also, three of them indicated to have a fair knowledge of logic, which may have helped in performing the second task, further reducing the total cognitive demands. In sum, their resources were sufficient to perform both tasks simultaneously.

There were two participants (3 and 5) that eventually had a strategy of being uninformative and a logical language use, but did not seem to use second-order ToM. These participants developed either their strategy or their language use during the experiment. These participants may represent an intermediate stage in which some resources are available for reflective cognition, but not yet enough to use all the complex skills. During the experiment the demands of the first tasks may have decreased because of automation, freeing more resources for the second task and thereby enabling the development of strategy or language use. These participants had enough resources available for doing the first task and part of the second task.

The third group consists of the six participants that did not show the use of any complex skills. It could be that these participants were too occupied with the first task of playing the game, to be able to also use reflective cognition. In other words, their resources may have been insufficient to perform both the first and the second task. These participants made relatively many mistakes in playing the game, so it could even be that the demands of the first task alone were too high already. This group may also include participants that did not think of logical language use as a possible way to use language. Some of the participants in this group developed strategies which they described as making things difficult for the opponent, but they did not relate this to the amount of information they revealed.

The use of quantifiers

Hypothesis 2 stated that in an uncooperative situation, people will shift their interpretation and production of quantifiers from pragmatic (according to Grice's quantity maxim) to less pragmatic (violating Grice's quantity maxim). None of the participants developed a more logical language use in the way that was meant in hypothesis 2. However, by the end of the experiment, half of the participants had a fairly to strict logical language use! As explained above, it could be the case that these participants had enough resources left for reflective cognition from the start of the experiment, so that they could use logical productions and interpretations. (Though our number of participants is very small, this percentage is comparable to and in between the percentages of adults giving logically correct responses to pragmatically infelicitous some sentences in (Noveck, 2001), and (Feeney et al., 2004).)

Automated versus deliberate processes

Hypothesis 3 stated that in interpreting and producing quantifiers, adults make use of an automated process, which results in pragmatic use of the quantifier, and that this process can be ‘overruled’ by a deliberate reasoning process, resulting in logical use of the quantifier. It seems that pragmatic language use is not automated for all adults in the situation of the experiment, since some participants developed pragmatic language use while repeatedly playing Master(s)Mind(s). This could be because of the context-dependent meaning of some of the scalar terms used.

By the end of the experiment, half of the participants had a non pragmatic language use. It is unlikely that they were unable to make pragmatic interpretations and productions, since they were all adults (being unable to draw pragmatic inferences would be a serious challenge in daily life). Thus, it seems that people can indeed choose to use language in a non pragmatic way. Since most participants wrote down comments on their way of interpreting the scalar terms, it seems that changing this interpretation is indeed a deliberate reasoning process.

Limited Theory of Mind use

It was found that although participants used first-order ToM, they did not use all kinds of reasoning possible and useful in this context with a first-order ToM. Most participants were aware of the desire of their opponent to know their secret code and wanted to make it as difficult as possible for the opponent to get that knowledge. Most of them considered how their opponent would interpret a certain feedback sentence. However, only few participants considered what information would be revealed, or would have to be revealed as a result of a certain guess they made themselves. Apparently, participants were more aware of information that was revealed by their language productions than they were of information that was revealed by their actions, or that would have to be revealed as a result of their actions.

Participants who used second-order ToM did use all the reasoning patterns possible with first-order ToM which were described above. Thus, they considered guesses as well as sentences to be informative.

5.3. GENERAL DISCUSSION AND CONCLUSIONS

Our experimental results provide some support for our dual-task model of skill acquisition. Though the evidence is not conclusive yet, it does present a problem for the theory of skill acquisition proposed by Fitts (Fitts, 1964). The performance of participants who were using second order Theory of Mind, and who were using language in a logical way, was fast and with relatively few errors. According to Fitts’ theory, this suggests that these skills result from automated processes. However, these participants were able to explicitly

write down their reasoning about their opponent and their thoughts about the use of quantifiers. In other words, they had full deliberate access to their skills, whereas Fitts predicted that deliberate access to automated skills is limited.

Our dual-task model is not challenged by the same problem. In our view skill acquisition is a continuous interplay between deliberate and automated processes. Because, for the advanced participants, performance was (partly) automated, they could apply reflective cognition to their advanced strategies for the game. Notice however that we do use Fitts' description of the process of automation as a building block for our dual-task model.

Fitts' theory predicts that the autonomous stage is reached by repeated performance of the task. Only a few participants in our experiment were able to benefit from repeatedly playing the game. Yet some participants who had never played this particular game before were able to use all complex skills from the start. Our dual-task model of skill acquisition gives a more elegant explanation for this than does Fitts' theory. These advanced players manage to benefit from other skills that they master, by applying them to the new context. In Fitts' theory however this is not possible, since deliberate access to such mastered skills is limited.

Our results corroborate well-known findings of Colman and others that most players easily manage first-order ToM, while many players are limited to first-order reasoning and almost all are limited to second-order reasoning (Colman, 2003; Hedden and Zhang, 2002). However, our research does not corroborate Hedden and Zhang's conclusions about skill acquisition (particularly second order ToM reasoning) during game playing (Hedden and Zhang, 2002). When Hedden and Zhang asked players to predict the other's moves in a game, it turned out that participants started out with the default assumption that their opponent would not model them. When in subsequent turns the opponent took a first-order strategy, the subject would switch to second-order predictions (but see (Flobbe et al., 2007; Colman, 2003) for criticism of Hedden and Zhang's methodology). Our findings do not show such an improvement of Theory of Mind use during the course of the repeated games. Instead, they are much more in line with the results of Flobbe et al., who in another context also find that their adult subjects start the experiment with proficiency in either first-order or second-order applied ToM, and do not move up the hierarchy of ToM when being exposed to the game for a number of different trials (Flobbe et al., 2007). Thus, the use of ToM, as well as the development of this skill were very limited in our study, even though most of our participants were advanced university students. Our results provide a warning against assuming that because some people (such as trained logicians) on some occasions apply higher-order Theory of Mind, such reasoning is at all widespread.

6. Future Work

In future work, more evidence for or against hypothesis 1 needs to be found. The difficulty of the first task –playing the game correctly– needs to be varied, in order to investigate whether this influences the use of skills that result from reflective cognition. In the Mastersminds-experiment, there are several ways to do so. The interface of the computer program used could be made less user-friendly, time pressure could be added, and the number of colors in a secret code could be varied. Also, a less well-known game similar to Mastersminds could be used, because most people know the regular version of the game Mastermind, which can be a benefit in playing Mastersminds. During the experiment, some participants got tired. In future experiments, fatigue could be measured by determining physical parameters, e.g. heart rate and blood pressure, in order to gauge to what extent advanced cognitive skills suffer from fatigue, which could be a measure for how much effort they require and thus how well they are mastered.

An improvement in the experimental setup should be made in order to better be able to measure complex skills and strategies. Participants with pragmatic language use had a disadvantage in strategy development. A strong strategy for this game is to reveal little information. The less informative sentences that logical language users could prefer often were regarded as false by pragmatic language users so that they could not use these sentences. By including more expressions, such as for example *niet alle (not all)*, the possibilities for pragmatic language users can be increased. In addition, logical language use should be taught to participants prior to the experiment, to make sure that all participants are aware of the possibility of using language in this way. Also, participants' motivation to use second order ToM could be increased by giving them a reward depending on the number of trials they need to correctly guess the opponents' secret code, and the number of trials the opponent needs to guess their secret code correctly. Finally, in the experiment, order of ToM was measured by the way participants described their strategies in the questionnaire. Unfortunately, talk-aloud protocols do not suffice for acquiring correct information about subjects' knowledge and reasoning strategies (Nisbett and Wilson, 1977). Thus, it would be good to also test order of ToM separately, for example by story tasks such as those used in (Flobbe et al., 2007).

An alternative for hypothesis 2 could be: *In an uncooperative conversation, some people will use quantifiers in a way that is not in accordance with Grice's quantity maxim, but that is more truth-conditional.* To test this hypothesis, it should be investigated whether the cooperativeness of the situation has an influence on language use. This could be done by observing the language use of the participants who had a logical language use during

the Mastersminds-experiment, while they play a fully cooperative game, in which a mutual goal has to be reached by two or more players.

To make it more clear whether or not logical language use can only result from overruling pragmatic language use, as stated in hypothesis 3, it would be interesting to let the participants to the Mastersminds experiment do an experiment like the one that was conducted by Feeney et al. on the interpretation of scalar terms in sentences like ‘Some fish can swim’ (Feeney et al., 2004). This could also be done for other scalar terms than *some*. Such an experiment could reveal whether the participants who had a logical language use from the start still need to overrule their pragmatic language use. If participants were to complete such an experiment before and after doing the Mastersminds-experiment, people who have shifted to more logical use are expected to have increased reaction times, since they now have to overrule their automated interpretation process.

Finally, our future project will be designing logical models that are less idealized than the usual epistemic logic S5 described in subsection 2.2 and that more accurately model different types of human reasoning on several orders of ToM. Thus far, there have been no logical models of human reasoning at higher orders of Theory of Mind. Stenning and Van Lambalgen have presented an interesting logical model of human reasoning in false-belief tasks that test for first-order ToM, based on closed-world reasoning as used in non-monotonic logics (Stenning and van Lambalgen, 2007). Van Ditmarsch and Labuschagne model different types of shortcomings in first-order ToM using frame-definable properties of different agents’ preference relations (where an agent prefers one state over another if he considers it more likely) (Van Ditmarsch and Labuschagne, 2007). For example, autists might think that ‘another agent’s preferences are exactly similar to my own’⁴.

7. Acknowledgments

We would like to express our gratitude to Petra Hendriks, who has been the insightful co-supervisor of Lisette Mol’s Master’s research project in 2004 and co-author of (Mol et al., 2005). We would also like to thank Niels Taatgen, co-author of (Mol et al., 2005).

Furthermore, Rineke Verbrugge gratefully acknowledges the NIAS (Netherlands Institute for Advanced Studies in the Humanities and Social Sciences) and NWO (the Netherlands Organization for Scientific Research, grants 051-04-120 and 400-05-710) for enabling her to work on this article as a NIAS fellow of the thematic group ‘Games, Action and Social Software’.

⁴ Stenning and Van Lambalgen warn against the common conflation of autism spectrum disorders with delays in correct first-order ToM (Stenning and van Lambalgen, 2007), see also Van Lambalgen’s article in this special issue.

Finally, we would like to thank the two anonymous referees for their useful criticisms and suggestions.

References

- Antoniou, G.: 1997, *Nonmonotonic Reasoning*. Cambridge (MA): MIT Press.
- Benz, A. and R. van Rooij: 2007, 'Optimal assertions, and what they implicate. A uniform game theoretic approach'. *Topoi* **26**(1), 63–78. special issue on logic and psychology, edited by J.F.A.K. van Benthem, H. Hodges, and W. Hodges.
- Binmore, K.: 1992, *Fun and Games: A Text on Game Theory*. Lexington (MA): D.C. Heath and Company.
- Blutner, R.: 2000, 'Some aspects of optimality in natural language interpretation'. *Journal of Semantics* **17**, 189–216.
- Colman, A.: 2003, 'Depth of strategic reasoning in games'. *Trends in Cognitive Sciences* **7**, 2–4.
- Dekker, P. and R. Van Rooij: 2000, 'Bi-directional optimality theory: An application of game theory'. *Journal of Semantics* **17**, 217–242.
- Dunin-Keppicz, B. and R. Verbrugge: 2006, 'Awareness as a vital ingredient of teamwork'. In: P. Stone and G. Weiss (eds.): *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'06)*. New York (NY), pp. 1017–1024, IEEE / ACM Press.
- Evans, J. S. B. T.: 2003, 'In two minds: Dual-process accounts of reasoning'. *Trends in Cognitive Sciences* **7**, 454–459.
- Fagin, R., J. Halpern, Y. Moses, and M. Vardi: 1995, *Reasoning about Knowledge*. Cambridge (MA): MIT Press.
- Feeney, A., S. Craffton, A. Duckworth, and S. Handley: 2004, 'The story of some: Everyday pragmatic inference by children and adults'. *Canadian Journal of Experimental Psychology* **58**, 90–101.
- Fitts, P. M.: 1964, 'Perceptual-motor skill learning'. In: *Categories of Human Learning*. New York: Academic Press.
- Flobbe, L., R. Verbrugge, P. Hendriks, and I. Krämer: to appear, 2007, 'Children's application of Theory of Mind in reasoning and language'. *Journal of Logic, Language and Information*.
- Grice, P.: 1989, *Studies in the Way of Words*. Cambridge (MA): Harvard University Press.
- Hedden, T. and J. Zhang: 2002, 'What do you think I think you think? Strategic reasoning in matrix games'. *Cognition* **85**, 1–36.
- Horton, W. and B. Keysar: 1996, 'When do speakers take into account common ground?'. *Cognition* **59**, 91–117.
- Johnson-Laird, P.: 1983, *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Cambridge: Cambridge University Press.
- Karmiloff-Smith, A.: 1992, *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge (MA): MIT / Bradford Books.
- Keysar, B., D. Barr, A. Balin, and T. Peak: 1998, 'Definite reference and mutual knowledge: Process models of common ground in comprehension'. *Journal of Memory and Language* **39**, 1–20.
- Keysar, B., D. Barr, J. Balin, and J. Brauner: 2000, 'Taking perspective in conversation: The role of mutual knowledge in comprehension'. *Psychological Science* **11**(1).
- Keysar, B., S. Lin, and D. Barr: 2003, 'Limits on theory of mind use in adults'. *Cognition* **89**, 25–41.

- Kooi, B.: 2000, 'Master(s)Mind(s)'. Technical report, Department of Philosophy, University of Groningen, Groningen. Manuscript at <http://www.philos.rug.nl/~barteld/>.
- Kooi, B.: 2005, 'Yet another Mastermind strategy'. *ICGA Journal* **28**, 13–20.
- McKelvey, R. and T. Palfrey: 1992, 'An experimental study of the centipede game'. *Econometrica* **60**, 803–836.
- Meyer, J.-J. C. and W. van der Hoek: 1995 / 2004, *Epistemic Logic for AI and Theoretical Computer Science*. Cambridge: Cambridge University Press.
- Mol, L.: 2004, 'Learning to Reason about other People's Minds'. Technical report, Institute of Artificial Intelligence, University of Groningen, Groningen. Master's thesis.
- Mol, L., N. Taatgen, R. Verbrugge, and P. Hendriks: 2005, 'Reflective cognition as secondary task'. In: B. Bara, L. Barsalou, and M. Bucciarelli. (eds.): *Proceedings of Twenty-seventh Annual Meeting of the Cognitive Science Society*. Mahwah (NJ), pp. 1925–1930, Erlbaum.
- Nagel, R. and F. Tang: 1998, 'An Experimental Study on the Centipede Game in Normal Form: An Investigation on Learning'. *Journal of Mathematical Psychology* **42**, 356–384.
- Nisbett, R. and T. Wilson: 1977, 'Telling more than we can know: Verbal reports on mental processes'. *Psychological Review* **84**, 231–259.
- Noveck, I.: 2001, 'When children are more logical than adults: Experimental investigations of scalar implicature'. *Cognition* **78**, 165–188.
- Papafragou, A. and J. Musolino: 2003, 'Scalar implicatures: Experiments at the semantics-pragmatics interface'. *Cognition* **86**, 253–282.
- Parikh, R.: 2003, 'Levels of knowledge, Games, and Group Action'. *Research in Economics* **57**, 267–281.
- Premack, D. and G. Woodruff: 1978, 'Does the chimpanzee have a theory of mind?'. *Behavioral and Brain Sciences* **4**, 515–526.
- Rosenthal, R.: 1981, 'Games of Perfect Information, Predatory Pricing, and the Chain Store'. *Journal of Economic Theory* **25**, 92–100.
- Schneider, W. and J. Chein: 2003, 'Controlled and automatic processing: Behavior, theory, and biological mechanisms'. *Cognitive Science* **27**, 525–559.
- Schneider, W. and R. Shiffrin: 1977, 'Controlled and automatic human information processing. I. Detection, search, and attention'. *Psychological Review* **84**, 1–66.
- Stenning, K. and M. van Lambalgen: to appear, 2007, *Human Reasoning and Cognitive Science*. Cambridge (MA): MIT Press.
- Sun, R. and X. Zhang: 2004, 'Top-down versus bottom-up learning in cognitive skill acquisition'. *Cognitive Systems Research* **5**, 63–89.
- Taatgen, N. A.: 2005, 'Modeling parallelization in skill acquisition'. *Cognitive Science* **29** (3), 421–455.
- Teepe, W.: 2006, 'Proving possession of arbitrary secrets while not giving them away: New protocols and a proof in GNY logic'. *Synthese (Knowledge, Rationality and Action)* **149**, 409–443.
- Van der Hoek, W. and R. Verbrugge: 2002, 'Epistemic logic: A survey'. In: *Game Theory and Applications*. New York: Nova Science Publishers.
- Van Ditmarsch, H. and W. Labuschagne: to appear, 2007, 'My preferences about your preferences: A case study in theory of mind and epistemic logic'. *Synthese (Knowledge, Rationality and Action)*.
- Van Eijck, J. and R. Verbrugge (eds.): to appear, 2008, *Discourses on Social Software*, Texts in Games and Logic. Amsterdam: Amsterdam University Press.
- Wason, P.: 1966, 'Reasoning'. In: B. Foss (ed.): *New Horizons in Psychology I*. Harmondsworth: Penguin Books, pp. 135–151.