THEORY OF MIND IN LIMITED BIDDING

Bachelor's Project Thesis

Claudia Frank, c.l.l.frank@student.rug.nl

Supervisors: dr. H.A. de Weerd, T. Buwalda, B. Arslan

Abstract: Humans have the ability to use theory of mind recursively. That is, they are able to explicitly attribute unobservable mental content to others and can engage in higher order theory of mind to consider what others believe about their own beliefs. Agent-based simulation studies of theory of mind show a clear advantage for the use of higher-order theory of mind in competitive games. In this paper we investigate whether humans use higher-order theory of mind to their advantage in the limited bidding game. Participants played several games against agents who use different orders of theory of mind. The results do not show that people used theory of mind in the limited bidding game but that people prefer to use simpler strategies instead.

1. Introduction

Humans are capable of theory of mind (ToM), that is the ability to reason about the unobservable mental states of the self and others (Premack & Woodruff, 1978). With theory of mind there is an understanding of others as mental beings like oneself, each with their own mental states, such as thoughts, wants, motives and beliefs. Theory of mind allows us to understand the behaviour of others, reason about the intention of others and predict the future behaviour of others.

Studies done on human reasoning in competitive games show that we use theory of mind recursively in decision-making processes (Perner & Wimmer., 1985; Meijering et al., 2011; Hedden & Zhang, 2002). This recursive ability is referred to as higher-order theory of mind and has been established both through false belief tasks (Perner & Wimmer, 1985) and strategic games (Hedden & Zhang, 2002; Meijering et al., 2011). Higher-order theory of mind can occur in different orders, such as a zero-order, first-order and second-order. To illustrate the differences between the orders, imagine a situation between the two cartoon characters Tom and Jerry (Barbera & Hanna, 1940). Every evening around 4 pm it is dinner time for Tom, the cat. Jerry, the mouse, observed that Tom repeatedly got fed at this hour everyday and predicts that this will happen again today. Here, Jerry is using a zero-order theory of mind because he is simply recalling a fact. Tom however, knows that Jerry knows that it is his dinner time. He is reasoning about the mental state of Jerry and is therefore using first-order theory of mind. Because he knows this, he predicts that Jerry will try to steal his food and decides to wait by his door to try to stop him. Jerry on the other hand, being the clever one in this cartoon series, knows that Tom knows that he knows that it is Tom´s dinner time and that he will try to steal his food and predicts that Tom would be waiting by his door to try to stop him. Jerry uses second-order theory of mind reasoning to infer this and decides to outsmart Tom by going to the cat bowl through another route. This example illustrates how theory of mind can be used in different orders and how it can be beneficial in predicting the opponents behaviour. Now, this example is one of two animal cartoon characters that are depicted with human-like behaviour. It has been proposed that some non-human animals also exhibit theory of mind (Tomasello, 2009; Parrish & Brosnan, 2012). Whether this behaviour can be considered equal to the theory of mind of humans has been a subject of debate (Penn & Povinelli, 2007). One major differences between human theory of mind and the theory of mind exhibited by other animals is that we can use theory of mind recursively. Children from the age of two to five already acquire full competence on first-order theory of mind tasks (Wellman et al., 2001). Adults have been shown to be capable of performing up to fourth-order theory of mind (Kinderman et al., 1998).

The benefits of higher-order theory of mind have been recently explored through agent-based modelling (de Weerd et al., 2013). Agent-based

modelling is a simulation technique in which individual agents act and interact based on their own perception of their local situation. Theory of mind was simulated in these agents through recursive modelling where the opponents are modelled as an opponent-modelling agent itself. The agents then played several competitive games against each other, including the limited bidding game. In the limited bidding game, the agents had to try and outbid each other through a course of five rounds using a limited set of bidding coins. Results from the simulation show a clear advantage for the agents that use a first-order and second-order theory of mind over the opponents that used a lower order theory of mind. For orders of theory of mind beyond the second the additional advantage was found to be marginal. These results were also consistent over the other competitive games studied in the agent-based simulation study (de Weerd et al., 2013).

By studying the differences between the ways in which human players play the limited bidding game against the different theory of mind agents, we explore the following research question: Do humans use a higher order theory of mind when playing the limited bidding game, and if so, to what extent? We hypothesize that humans would either use a first or second-order theory of mind in the limited bidding game, because it is shown to be most beneficial, and would be able to switch to lower order theory of mind according to the perceived benefit of using such order. This means that if the player is capable of using a second-order theory of mind but notices that the opponent is a zero-order, he could be able to think in a first-order theory of mind manner in order to better predict the action of the opponent.

2. Method

2.1. Participants

Data has been collected from nineteen participants. The participants were all college students (Mean age =22, SD=1.9) and have been recruited from a paid research participants group on facebook. There were in total twelve females and seven males. All participants gave informed consent prior to admission into the study and were properly debriefed about the nature of their opponents at the conclusion of the experiment. The experiment lasted about one hour. As a compensation for participating in this study, the participants were paid eight euros.

2.2. Apparatus

The limited bidding game was implemented in JavaScript and was presented to the participants via a computer screen. The collected data were analyzed in R.

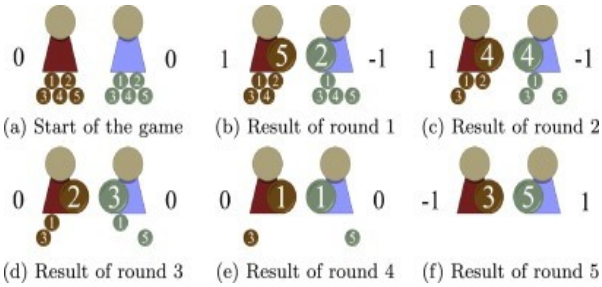

Figure 2.1: The limited bidding game reprinted with permission from de Weerd et al. (2013).

The limited bidding game is an adaptation from a game presented in De Bono (1998). Figure 2.1 shows how the limited bidding game is played among the two players. Each player is handed at the beginning of the game five coins, from value one to five. Over the course of five rounds, the players simultaneously choose one of their coins to bid against the other for that round. Once the players choices have been made, the coins played are revealed and compared. The player with the highest bid wins the round. In case of a draw, there are no winners. For each round won, the player will win a point and for each round lost, the player will lose a point. The objective of the game is to win as many rounds as possible while losing as few rounds as possible. Each coin can only be used once in each game. The players would therefore have to weigh the additional probability that they will win the current round with a higher valued token against the loss of competitive strength in later rounds. It is not possible to win all rounds. Instead, each player can win a maximum of four rounds. As a result a player can achieve a maximum score of three in the limited bidding game.

2.3 Opponents

Each opponent was simulated by an agent that either uses theory of mind or a randomizing strategy. The agents that use theory of mind have been implemented accordingly to the agent-based model used in a simulation study of limited bidding (de Weerd et al., 2013). The

order of theory of mind ranged from a zero-order to a second-order. The opponents were presented together with a designated name. The names were to encourage the participants to figure out the strategy used by that agent without being affected by the actions of the previous agent.

The following orders of theory of mind were presented to each participant:

Zero-order theory of mind

The zero-order theory of mind (ToM0) agent is unable to model the mental content, such as beliefs and intentions of others. This agent therefore does not really use theory of mind and simply relies on its memory of past events to predict the opponents future actions. In the limited bidding game it could be for example that the opponent would always start the game by bidding with a four. An agent that uses a zero-order theory of mind would observe this and bid with a five at the start of the next game. This agent is intended to model an inexperienced player who believes that the opponents actions can be predicted solely by what happened in the previous rounds, without considering the fact that the opponent might be making choices based on the agent´s actions.

First-order theory of mind

The first-order theory of mind agent (ToM1) considers the possibility that the opponent is trying to win the game by reacting to the choices made by the agent. To predict the behavior of the opponent, this agent would place itself in the position of the opponent and consider the information that is available to them from their perspective. The first-order theory of mind agent therefore thinks one step further than the zero-order theory of mind agent would.
In the case of the limited bidding game, this would mean that the agent would consider what the human player could have observed from the agent. It could have been for example that the agent often started the game by bidding with a four. Considering that the human player observed this and would possibly start the next game with bidding a five, the agent would bid with a one instead at the beginning of the next game to get rid of his lowest coin.
The first-order theory of mind agent models his opponent as being able to use a zero-order theory of mind, but does not know the extent of the abilities of their opponents with certainty. Through repeated interaction, this agent may come to the belief that his opponent is not a zero-order theory of mind agent and that they do not have any beliefs at all. The agent could then adjust its order of theory of mind and play as a zero-order theory of mind agent to better fit the situation present.

Second-order theory of mind

The second-order theory of mind (ToM2) agent models his opponent as being able to use a first-order theory of mind. That is, a second-order theory of mind agent considers the possibility that his opponent is putting themselves in their position, and models him as a first-order theory of mind agent. In the limited bidding example where the agent would always start the game with a four and would now bid with a one based on first-order theory of mind reasoning. The agent would think a step further and consider the possibility that his opponent is using a first-order theory of mind. He would reason that his opponent predicts that he will bid with a one. He would predict that the opponent will therefore bid with a two and decides to bid with three instead.

Random agent

Additional to the agents that simulate different orders of theory of mind, the players played against an opponent that choses a random coin for each round during the game. The results obtained from playing against the Random agent will be used as a means of comparison for the results from the theory of mind agents. Because it is impossible to predict the moves of the Random agent, it is expected that the participants do not significantly increase or decrease in scores when playing against this opponent.

2.4 Procedure

Upon arrival, demographical data were collected from the participants including age and gender. The participants received instructions on how to play the game and the goal of the game. All instructions were presented via a computer screen. The participants were notified about the use of different strategies among the opponents but were not told what strategies were used. All participants did a test trial of one game before the start of the experiment, to get acquainted with the buttons. The experiment consisted of

playing 30 games against each opponent. This led to a total of 120 limited bidding games being played, which lasted about one hour. The scores of both players and the name of the opponent were visible on the screen. The scores of the players do not reset to zero when the player switches from opponents.
The participants were asked to make a prediction prior to making a decision on the bidding coin for each round. The participants were allowed to click on a ´no prediction´ button for when they could not predict the opponents behaviour.
After playing against all opponents the participant filled out an evaluation form involving questions on the game and theory of mind usage during game play.

2.5 Design

In order to assess what order of theory of mind is used among the participants in the limited bidding a within subject design was used. The participant played against each opponent for one block, which consisted of 30 consecutive games. The order and the designated names of the four opponents were randomly set for every participant. In every limited bidding game there are five rounds, for which the chosen coin, the prediction of the human player and the agent were recorded by the computer. Additional to the choices and predictions made, the reaction time of the human player was also recorded. In total 2850 data points were collected per opponent condition.

2.5 Analysis

To research the use of higher order theory of mind in the limited bidding game, the bidding choices made by the participants were analyzed using a random-effect bayesian model selection (Daunizeau et al., 2009). In the random-effect bayesian model selection (RFX-BMS), each strategy is a model with one free parameter lambda. This parameter can be different for each participant. In this analysis, we looked at 11 different models of participant behavior. RFX-BMS assumes that these are the only models that exists in the population and that the models have a fixed unknown distribution among the population. It is therefore important that as many models as possible are used for the analysis. It is also assumed that the subjects are sampled from a homogenous population with one unknown mode.

For every participant and strategy the RFX-BMS will calculate the best fitting model. The model fits of the strategies across all participants are then compared, in order to decide the most likely way that the strategies are distributed across the participants.
The following strategies were used as models in the RFX-BMS:

Drift+x
When the Drift+x strategy is used, the agent would make x choice higher than they did last time they were in the same situation with probability lambda and would randomly choose a different action with probability (1- lambda). Both Drift+1 and Drift+2 are models used in the analysis.

Sticky
The Sticky strategy could be regarded as a Drift+0 strategy. When the agent is in the same situation as it was before, it would repeat the choice it did last time with probability lambda. It would randomly choose a different action with probability (1-lambda).

Win-shift-lose-shift (WSLS)
An agent that uses WSLS strategy, like Sticky, would repeat the last choice it made in a similar situation with probability lambda only if it won with that choice last time. It would randomly pick another action if it lost or tied last time it was in that situation.

Bias
The Bias strategy chooses the lowest coin with some probability lambda, and randomized over the rest of the coins with probability (1-lambda).

Random
The agent would randomly pick an action in each situation.

Theory of mind
The theory of mind models make choices in the same manner as the theory of mind agents that the participants play against. In the RFX-BMS, ToM0 through ToM4 are used as models.

The RFX-BMS would output how closely the participant data fits with a certain model in comparison with the other models. To test if the results are significantly different among the four

opponents during the game a Chi-Square test for independence is used.

An analysis of variance (repeated measures ANOVA) was conducted to indicate if there are significant differences between the win scores, correct prediction scores, no prediction scores or reaction time for the different opponents. A significant decrease in score or increase in reaction time for a higher-order theory of mind opponent would suggest the order of theory of mind used among the participants.
To research if fatigue had an effect on the participants performance, an analysis of variance was also done to indicate differences among the different rounds. All statistical analysis were conducted in Rstudio (R Core Team, 2013).

3. Results

Data was collected for all 19 participants for which each participant completed 120 games of limited bidding. The participants played 30 games against each opponent. Per game, five predictions and decisions were recorded from the player, along with the reaction time and the decision made by the computer agent. The number of correct predictions and the number of times the human player won for each participant were calculated from this data. Furthermore, the scores obtained after playing against each opponent were calculated by subtracting the amount of times the player lost from the amount of times they won during those 30 games against the opponent.

3.1. Performance

Participants could choose to not make a prediction when they could not predict what the opponent was going to bid with by clicking on a ´no prediction´ button before choosing their bidding coin. Participants that excessively used the no prediction button(>100 times for a block) were considered as outliers and were excluded from the analysis of variance for correct prediction.
The amount of times the participant correctly predicted the opponents move and the number of times the participant won against the agent were compared for the different opponents (See Figure 3.1). It can be observed in Figure 3.1 that there is a slight increase in prediction and winning scores in the case of an increased order of theory of mind opponent.

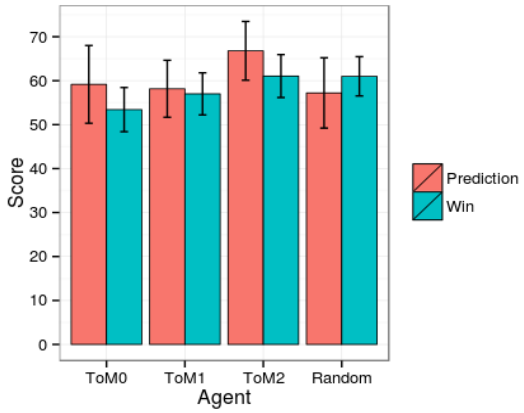

Figure 3.1: Prediction and win scores for the different opponents with 95% confidence interval.

An analysis of variance did not show significant effect of agent strategy or theory of mind level of the opponent on the prediction scores $F(3, 56) = 1.403$, $p>0.05$ or on the winning scores $F(3, 68) = 2.468$, $p>0.05$.
The mean prediction score (M=66.27, SD=10.65) is close to the mean score if they were to randomly select a prediction each turn (M=68.5). This score is calculated by multiplying the chance that they would get the prediction correct for a round with the amount of points possible in that round during the whole block. A mean prediction score close to this score suggests that the participants were unable to reliably predict the decision of the opponent.

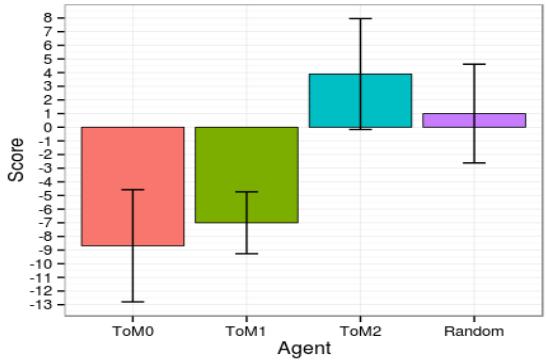

Figure 3.2: Scores after playing against the different opponents with 95% confidence interval.

The difference in scores obtained after playing against the opponents can be seen in Figure 3.2. It can be observed that participants won when playing with a second-order theory of mind agent and lost mostly when playing against a first-order or zero-order theory of mind agent.

As expected playing against the Random agent resulted in neither losing or winning.
An analysis of variance on these scores yielded significant variation among the opponent conditions, $F(3, 68) = 3.951$, $p<0.05$. A post hoc Bonferonni test showed that the second-order theory of mind scores and the other theory of mind scores differed significantly at $p<0.005$. There is also a slight increase in score observable for when the participants played against a first-order theory of mind agent compared to when they played against a zero-order theory of mind agent. These results suggest that the participants found opponents with a lower order theory of mind agent more difficult to defeat than opponents that use higher order theory of mind.
An analysis of variance on the reaction time to make a prediction and bidding decision did not show any significant difference for for the different opponents $F(3, 68) = 0.125$, $p>0.05$.

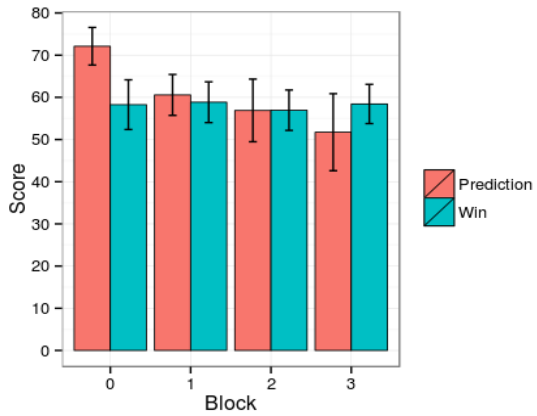

Figure 3.3: Prediction and win scores for the different blocks with 95% confidence interval.

To research if there was any effect of fatigue, differences in the use of the ´no prediction´ button, the correct prediction scores, win scores, the overall scores and the reaction times were also compared for the different blocks. A block consisted of 30 games against one opponent.
Figure 3.3 indicates that there is a decrease in the amount of correct predictions made for later blocks. An analysis of variance did not show any significant effect of the block on the winning scores $F(3, 68) = 0.098$, $p>0.05$ or the correct prediction scores $F(3, 56) = 0.767$, $p>0.05$.
The block appears to have an effect on the ´no prediction´ button use (See Figure 3.4), the scores (Figure 3.5) and the reaction time to make a prediction and choose a bidding coin (See Figure 3.6). As the block increases, the participants seem to have made more use of the no prediction button, worsened in score and took less time to make a choice for the prediction coin and bidding coin.

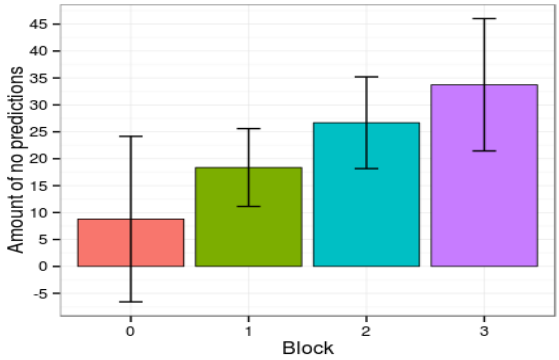

Figure 3.4: Amount of no predictions made for the different blocks with 95% confidence interval.

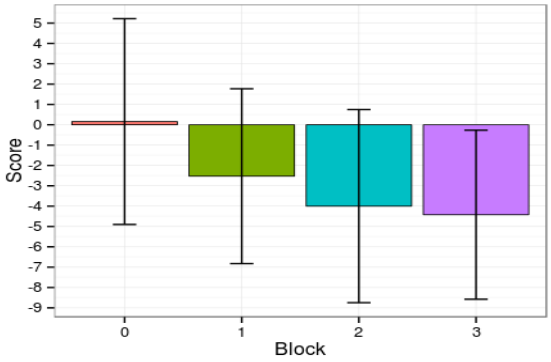

Figure 3.5: Scores for the different blocks with 95% confidence interval.

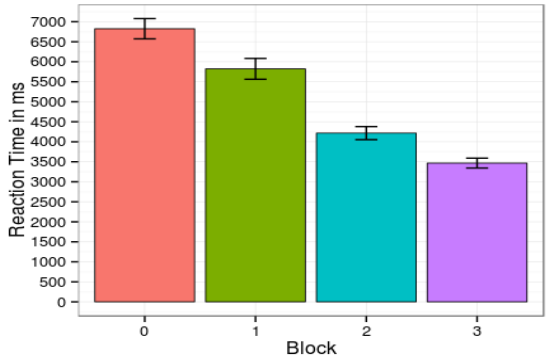

Figure 3.6: Average reaction time for the different blocks with 95% confidence interval.

However, an analysis of variance did not show a significant difference in scores $F(3, 68) = 1.234$, $p>0.05$, reaction time $F(3, 68) = 2.52$, $p>0.05$ and in amount of no predictions made for the different blocks $F(3, 68) = 0.345$, $p>0.05$. It is important to take in consideration that the

results found for the differences between the blocks could have been affected by the random distribution of the opponent strategy across the different blocks.

3.2. Strategies used

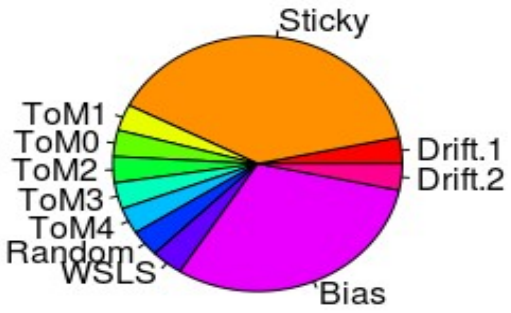

Figure 3.7: Proportions of strategies used in limited bidding.

The RFX-BMS analysis on the choices made by the agent and player produced the following distribution of strategies that were used by the participants during the limited bidding game (See Figure 3.7). Both Sticky and Bias obtained remarkably high percentages compared to any of the theory of mind strategies used in the analysis. This shows that the participants prefer to use simple strategies, such as making the same bidding decision as they did last time when they were in a similar situation (Sticky) or always bidding with the lowest coins (Bias). The Sticky strategy obtained a relatively higher percentage than the WSLS strategy. This could indicate that the participants did not learn from previous mistakes made in similar situations and would prefer to repeat their last action rather than choosing a different one. The obtained percentages for the theory of mind strategies were similar to the obtained percentages for the random model. This suggests that the probability of the participants choosing a random bidding coin in each turn is as high as the probability for them to be using a form of theory of mind when playing the limited bidding game. To research if the participants switch strategies depending on their opponent, the probabilities of the strategies were also calculated for the different opponents the players played against. Figure 3.8 depicts the strategies that showed difference among the different opponents. The strategies Bias and Sticky remain for the individual opponents with remarkably high percentages compared to the other strategies, including the theory of mind strategies. Although Figure 3.8 indicates some differences in strategies, a test for independence shows that the percentage of participants that use a certain strategy did not differ for the different opponents X^2 (30) = 4.25, p> 0.05.

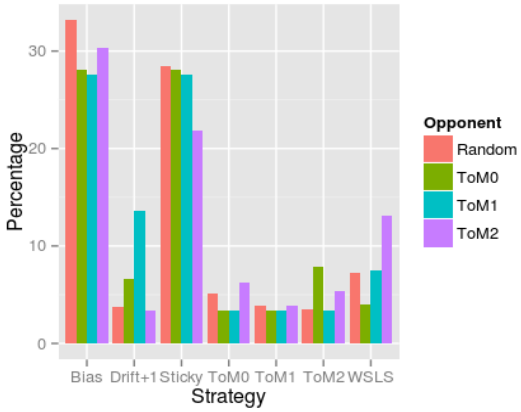

Figure 3.8: Estimated percentages of the strategies used scores across the different opponents.

To determine if there were differences between strategies used at the beginning of a game versus at the end of a game, the probabilities were calculated for the first three and last three rounds in a game separately (See Figure 3.9). It is observed that there are differences in the percentages between the Bias, Drift+3 and Sticky strategies but not any differences in percentages between the theory of mind strategies. In particular, Sticky seems to have a strong decrease in probability for the last three rounds compared to the first three. For the other simple strategies used in the analysis, no differences were found and were therefore not depicted in Figure 3.9. A test for independence shows that these observed differences are not significant and that the percentage of participants that use a certain strategy did not differ significantly between the first three rounds and the last three rounds $X^2(12) = 6.29$, p>0.05.

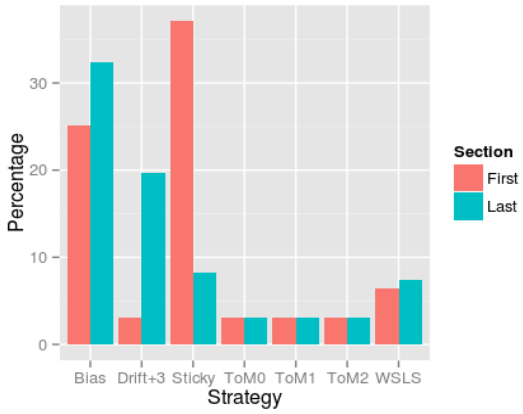

Figure 3.9: Estimated percentages of the strategies used scores across the first 3 rounds and the last 3 rounds.

3.3 Evaluation forms

After playing the limited bidding game against the four opponents, the participants were asked to fill in an evaluation form concerning the experiment. The form consisted of questions involving fatigue, difficulty and theory of mind use. Eight out of nineteen participants reported feeling fatigue from the experiment. Almost all participants noticed that the agents were using different strategies and found it difficult to defeat them. Twelve out of nineteen participants felt motivated enough to play attentively during the whole experiment. Loss of motivation was reported due to not being able to find a clear pattern in their opponent´s moves or losing hope of winning the game due to many losses.

Besides indication of fatigue or loss of motivation, answers gathered from evaluation papers also point out to theory of mind use during the game. Eight out of nineteen participants reported looking for patterns in the opponents behaviour, which is considered a zero-order theory of mind. Five participants reported using a first-order theory of mind where they would place themselves in the opponent´s position and use that information to try to outsmart them. Two participants reported using a specific strategy, which consisted of intentionally losing to the agent in the beginning in order to stay with the high valued coins later in the game. This resembles the Bias strategy, where the player would always prefer to bid the lower bidding coins.

4. Discussion

In this study the recursive use of theory of mind has been explored in the limited bidding game. Studies done on agent-based modelling of theory of mind show that it is beneficial to use a first-order theory of mind and second-order theory of mind over the opponents that use lower orders theory of mind. In this study we asked the research question: ´Do humans use a higher-order theory of mind in the limited bidding game and if so to what extent?'. We hypothesized that humans would use a first-order or second-order theory of mind to make decisions in the limited bidding game and that they would be able to switch between these orders depending on the opponent.

To answer the research question we let nineteen college students play the limited bidding game against the theory of mind agents. The participants played against a zero-order, first-order and second-order theory of mind agent. Besides the theory of mind opponents the participants also played against an opponent who bids randomly. In every round the participants decided on which coin they would bid with and a prediction of what they would think the opponent will bid with. The option was also present to not give in a prediction.

The results suggest that the participants were more likely to use simpler strategies than theory of mind. The decisions made by the participants resembled more the decisions made by someone who would use a Sticky or Bias strategy. That is, the person would either simply repeat decisions that were made in the past for similar situations or would always choose the lowest coins to bid with. Regardless of the outcome of previous actions done in similar situations, the participants would repeat the action. This suggests that participants do not learn from previous mistakes made in the game. The strategy to bid with the lowest coin first was reported as a strategy used by the participants in the evaluation forms. A comparison between the strategies used for the first three rounds and the last three rounds did not show a significant difference in strategies. It is important to keep in consideration that these simple strategies were the ones that fit the decisions made by the participants best out of the ones used in the RFX-BMS. It could however be that the strategy used by the participants was not one that was

modelled in the RFX-BMS. What is evident from the results is that the simpler strategies were preferred over theory of mind in the limited bidding game. The use of theory of mind has been shown to be less than optimal in complex social situations (Flobbe et al., 2008; Keysar et al., 2003; Hedden & Zhang, 2002; Goodie et al., 2010). A reason often proposed for the lack of theory of mind found is that theory of mind can be difficult and cognitively demanding (e.g Verbrugge & Mol, 2008). It is possible that simpler strategies were preferred over theory of mind in the limited bidding game because they require less mental effort. Despite reports from the evaluation form indicating that the participants were aware of differences in strategies among the opponents, the results do not show significant differences in strategies used by the participants against different opponents. The participants still preferred simpler strategies over theory of mind for the different opponents and did not adjust their strategies according to the opponent´s strategy. It could be that the participants did not adjust their strategies, because they simply found these strategies to be the most efficient in all conditions. It could also be that no differences were found in strategies, because the strategy adjusted to was not modelled in RFX-BMS.

The number of times the participants won against the opponent, the number of times the participant made a correct prediction of the opponent´s actions and the reaction times of the participants to make this prediction and bidding decision were unaffected by the strategy used by the opponent. A significant difference however was found for the overall scores between the opponents. This score was present on the computer screen for every round during the 30 games and is, at the end of the block, equal to the difference between the amount of times they won and the amount of times they lost in a block. The participants scored lower when playing against a zero-order theory of mind agent, as opposed to playing against a first-order theory of mind agent. The score was significantly improved when playing against a second-order theory of mind agent. It was primarily expected that there would be an advantage when playing against lower-order theory of mind agents and thus an increase in score. The results indicate to our surprise that a lower-order theory of mind opponent was more difficult to predict and play against. The significant increase in score when playing against the second-order theory of mind agent could be due to the fact that the agent always played three in the first round of every game. This secured at least one point in prediction and in win scores for the participant in every game. The reason for the second-order theory of mind opponent to always start with a three is unclear.
Although the increase in score when playing against the first-order theory of mind compared to playing against a zero-order theory of mind was not significant, the observed difference could be explained through the use of simple strategies. Because there is indication that the participants often repeated their actions from previous similar situations and the zero-order theory of mind agent bases its belief of the opponent solely on what was done it the past, it can be argued that the zero-order theory of mind agent would be better able to predict the decisions of the participants than the higher-order theory of mind agents. The higher-order agent would first consider the opponent as being able to use theory of mind. Considering that theory of mind was not used by the participants, the higher-order theory of mind opponents would already lose some points by over-thinking about their opponents behaviour before adjusting their theory of mind level to a zero-order theory of mind.

Both the results obtained from the RFX-BMS analysis and results obtained from the analysis of variance of scores between the different opponents do not support our hypothesis that humans use a first or second-order theory of mind in the limited bidding game and have the capability to switch between orders of theory of mind according to perceived benefit of using such order. If the participants were using first-order theory of mind when playing the limited bidding game a decrease in the amount of correct predictions made and score when playing against a second order theory of mind agent would be expected. This was not the case for both the prediction scores, win scores and overall scores. If the participants use a second order theory of mind when playing the game it is still expected that they would adjust order of theory of mind between opponents and increase in score when playing against a lower order theory of mind agent. Furthermore, analysis on the bidding decisions made during the game do

not show that the participants use theory of mind or switch strategies among opponents.
However, results from the evaluation papers indicate that the participants do use theory of mind. Almost half of the participants reported using a zero-order theory of mind, five participants reported using a first-order theory of mind. The use of specific strategies were also reported in the evaluation papers, such as intentionally losing to the agent in the beginning in order to stay with coins of higher value later in the game.
The possible explanation is thus proposed for the higher probabilities found for the simple strategies compared with the theory of mind strategies and the conflicting results found in the evaluation papers concerning theory of mind: humans use theory of mind in a less strict manner than how it is implemented in the theory of mind agents. Instead of investing a lot of mental energy in applying theory of mind for every decision made in the limited bidding game, humans would use simpler strategies repeatedly over the course of a few games before considering the mental state of their opponent.

Research limitations

In order to asses if the data is reliable enough to measure theory of mind, the following question should be addressed: 'Were the participants motivated enough to play the game with full effort in order to defeat the opponent?'. There are indications that this was not the case. Firstly, the participant´s average prediction score was close to the average prediction score if the participants were to randomly select a prediction every turn in a game. This indicates that the participants were unable to reliably predict the opponents behaviour, or were not putting in full effort to predict the opponent´s actions. Secondly, the participants seem to have made more use of the ´no prediction´ button, worsened in score and took less time to make a prediction and decision on bidding coin as the blocks progressed. This could indicate that the participants started making more hasty decisions in later blocks and therefore affected their results. Thirdly, it was reported in the evaluation papers that there was loss of motivation and feelings of fatigue when during the experiment. A decrease in motivation and fatigue during the experiment might have thus affected the results obtained. Furthermore, the results obtained from the RFX-BMS analysis are limited to the strategy models used in the analysis.

5. Conclusion

In this paper, higher-order theory of mind use was explored in the limited bidding game. Regardless of the use of higher-order theory of mind being beneficial in the limited bidding game, we found that people prefer to use simpler strategies instead of theory of mind in their decision-making processes. A reason to use simpler strategies over theory of mind is that they are less cognitively demanding. Despite these results, theory of mind was still reported in the evaluation forms. Future research could investigate if theory of mind is still used but in combination with simpler strategies.

References

Barbera, J., & Hanna, W. (Writer and Director). (1940). *Tom & Jerry*. [Television programme].United States: Metro-Goldwyn-Mayer.

de Bono, E. (1998). *Edward de Bono's super mind pack: Expand your thinking powers with strategic games & mental exercises.*London: Dorling Kindersley.

Daunizeau, J., Friston, K.J., & Kiebel, S.J. (2009). *Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models* Physica D: nonlinear phenomena 238: 2089-2118.

Flobbe, L., Verbrugge, R., Hendriks, P., & Krämer, I. (2008). Children's Application of Theory of Mind in Reasoning and Language. *Journal of Logic, Language, and Information*, 17(4), 417-442.

Goodie, A. S., Doshi, P., & Young, D. L. (2012). Levels of theory-of-mind reasoning in competitive games. *Journal of Behavioral Decision Making*, 25:95–108.

Hedden, T., & Zhang, J. (2002). What do you think I think you think?: Strategic reasoning in matrix games, *Cognition*, 85(1), 1-36.

Keysar, B., Lin, S., & Barr, D. J. (2003). Limits on theory of mind use in adults. *Cognition*, 89(1), 25-41.

Kinderman, P., Dunbar, R. I. M., & Bentall, R. P. (1998). Theory-of-mind deficits and causal attributions. *British Journal of Psychology*, 89, 191–204.

Meijering, B., van Rijn, H., Taatgen, N., & Verbrugge, R. (2011). I do know what you think I think: Second-order theory of mind in strategic games is not that difficult, Proc. 33rd Annu. Conf. Cogn. Sci. Soc., 2486–2491.

Parrish, A.E., & Brosnan, S.F. (2012). Primate Cognition, In *Encyclopedia of Human Behavior (Second Edition)*, edited by V.S. Ramachandran, Academic Press, San Diego, 174-180.

Penn, D.C., & Povinelli, D.J. (2007). On the lack of evidence that non-human animals possess anything remotely resembling a 'theory of mind'. *Philosophical Transactions Of The Royal Society Of London. Series B, Biological Sciences*,*362*(1480), 731-44.

Perner, J., & Wimmer, H. (1985). "John thinks that Mary thinks that....": Attribution of second-order beliefs by 5- to 10-year-old children. *Journal of Experimental Child Psychology*, 39, 437-47.

Premack, D., & Woodruff, G. (1978) Does the chimpanzee have a theory of mind?, *Behavioral Brain Science*, 515-526.

R Core Team (2013). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria.

Tomasello, M. (2009). Why we Cooperate. Cambridge, MA: MIT Press.

Verbrugge, R., & Mol, L. (2008). Learning to Apply Theory of Mind. *Journal of Logic, Language, and Information*, 17(4), 489-511.

de Weerd, H.A., Verbrugge, R., & Verheij, B. (2013). How much does it help to know what she knows you know? An agent-based simulation study, *Artificial Intelligence*, Volumes 199–200.

Wellman, H.M, Cross, D., & Watson J. (2001). Meta-analysis of theory-of-mind development: the truth about false belief. *Child Development*,*72*(3), 655-84.