Effects of Feedback on Human Performance in Negotiation with a Metacognitive ACT-R Agent

(Bachelor's Thesis)

Tom Renkema, T.R.S.Renkema@student.rug.nl,
Supervisors: N.A. Taatgen*, C.A. Stevens*

July 5, 2015

Abstract

In this study, an experiment was conducted in which human subjects played a mixed-motive negotiation game, called Game of Nines, against a metacognitive ACT-R agent. In this game, two players had to divide up 9 points between them. The goal for both players was to get as many points as possible in order to get a higher score. The aim of this study was to examine whether participants' performance on Game of Nines could be enhanced as a result of receiving feedback. This feedback consisted of the strategy the agent was currently using and how it responded to the player's actions. The results showed an improvement over time across all participants, but no overall effect of feedback was found. However, splitting up participants into medians sorted on their obtained score did show a positive effect of feedback for the best performing subjects. These participants showed a large improvement in score over time, whereas the best performing subjects who did not receive feedback did not show any further improvement. This suggests that people do improve in negotiation by training against a metacognitive agent, but that only the best performing players in Game of Nines benefit from the feedback and adapt their own strategy to the agent's strategy. This leads to a higher score and thus to a better performance in negotiation.

*University of Groningen, Department of Artificial Intelligence

1 Introduction

Negotiation is an important social interaction. Being a good negotiator has benefits in many fields, ranging from social conflicts to international politics (Pruitt and Carnevale (1993) ; Adler (1986)). Several studies have examined if people can be trained to become better negotiators (e.g. Lin et al. (2014) ; Susskind and Corburn (2000)). This study examined whether people could improve in negotiation by playing against a metacognitive agent.

An experiment was conducted in which human subjects played a negotiation game called Game of Nines (Kelley et al., 1967). In this game, 9 points need to be divided between two players in each round. Both players can make offers to each other about the amount of points they would like to receive in a turn-based fashion. The round ends when either an agreement has been made or one of the players has quit. In the case of an agreement, both players receive the amount of points they agreed on, but neither receives any points if a player quits. The goal for both players is to get as many points as possible in order to get a higher score.

The negotiation scenario of Game of Nines is a distributive negotiation, because a fixed amount of value needs to be divided (Bartos, 1995). The negotiation is also a mixed-motive situation; both players want to gain as many points as possible, but are also forced to make an agreement with their opponent, because they receive no points at all if they fail to do so.

A negotiator has the option using different strategies in a negotiation scenario. Although there are many ways one could negotiate, two distinct strategies are commonly distinguished in literature: cooperative and aggressive. A cooperative strategy aims to maximize agreement between parties, makes concessions more often and typically makes more moderate opening offers (Komorita and Esser (1975) ; Chertkoff and Conley (1967)). In contrast, an aggressive strategy characteristically aims to maximize its individual gain, makes higher opening offers and makes fewer concessions (Hüffmeier et al., 2014).

Furthermore, one strategy might not be suitable against another strategy. For example, using a cooperative strategy against an opponent who uses an aggressive strategy risks being exploited (Hüffmeier et al., 2014). Therefore it is important for a negotiator to adapt their own strategy to their opponent's strategy. This means one should think about what their opponent thinks. This is referred to as metacognition (Metcalfe and Shimamura, 1994).

In a recent study conducted by Stevens (2015), participants played Game of Nines against two agents. One agent used a fixed aggressive strategy and the other a cooperative strategy. The results suggested that only a quarter of the participants adapted their strategy to the agent's strategy. The participants who adapted their strategy to that of the agent performed better than those who did not, as evidenced by their higher scores. This indicates that most participants did not use metacognition when negotiating with the agent.

The aim of this study was to examine whether participants' performance on Game of Nines could be enhanced as a result of receiving feedback on the agent's current strategy. This feedback consisted of the strategy the agent was currently using and how it responded to the player's actions. This means that the participants received information about how their opponent thinks, which does not occur in real life negotiations between humans. The feedback was intended to encourage participants to adapt their own strategy to that of the agent.

An experiment was conducted in which participants played Game of Nines against a metacognitive agent. This agent could switch between an aggressive and a cooperative strategy. The agent determined which strategy to use by imitating the strategy of its opponent.

It was hypothesized that participants would make use of the feedback by adapting their strategy to the metacognitive agent's strategy and obtain a higher score by doing so.

2 Methods

2.1 Game of Nines

Game of Nines is a mixed-motive negotiation game first used in a study conducted by Kelley et al. (1967). In this game, two players had to divide up points between them. The goal for both players was to get as many points as possible in order to get a higher score. The game consisted of multiple rounds in which 9 points were to be divided in each round. Both players could make offers to each other in a turn-based fashion in order to set an agreement about the division of points. The round ended if an agreement had been made or if one of the players had quit. Both players received the amount of points they agreed on in the case of an agreement, but neither received or lost any points when a player had quit. Each player also had the opportunity to make a final offer. In that case, the opponent had to either accept or reject the offer and could not make a counteroffer.

Additionally, both players received a Minimum Necessary Share value (MNS value) in each round. The MNS value is the minimum amount of points the player needed to receive in order to add points to their score. This was referred to as profit. For instance, if a player received 5 points while having a MNS value of 2, the player's profit was 3 points. A player could also have a negative profit if they accepted an offer which was lower than their own MNS value. Each player's MNS value was only presented to themselves, not to their opponent.

2.2 Metacognitive agent

In this study, a metacognitive agent served as an opponent for the human player. This agent was designed in a study by Stevens (2015) and was developed in the cognitive architecture ACT-R (Anderson et al., 2004) .

The agent distinguished between two negotiation strategies: aggressive and cooperative. The cooperative strategy aimed to maximize agreement between the human player and the agent. This means it made more moderate opening offers and was more likely to concede to the other player. The aggressive strategy on the other hand made higher opening bids and conceded less often.

The metacognitive agent aimed to use the same strategy as its opponent. This was done using instance-based learning (Aha et al., 1991). Multiple instances of possible scenarios were entered in the agent's declarative memory. Each instance consisted of multiple variables which described the context in which a move occurred. Instances were labeled with its corresponding strategy. The agent analyzed the opponent's offer and matched it with its instances in declarative memory. The strategy found in the matched instance determined the strategy it would use for its own counteroffer. See Stevens (2015) for an extensive explanation of the mechanics behind the agent.

2.2.1 Interface

A human player played against the metacognitive agent using the interface pictured in Figure 1.

The player could interact with the interface using the mouse. The player could submit an offer by checking one of the radio buttons labelled 1 till 9 and subsequently press the 'Submit Offer' button. To submit a final offer, the player had to check the corresponding box before submitting it. A participant could accept the agent's offer or quit the round by pressing the matching button. The player's current MNS value and score for the current block was shown on the left at all times. The number of bids occurred in a round was also shown. This was provided to give the player additional confirmation that their offer had been submitted and at what stage of the round the player was in. The model communicated through the chat window above. Whenever the agent made an offer, it also showed the amount of points the player would receive if they were to accept that offer. This was done in order to prevent the player focusing on the arithmetic too much, which could distract from the actual negotiation purpose of the game. The player still had to subtract the MNS value to calculate their potential profit.

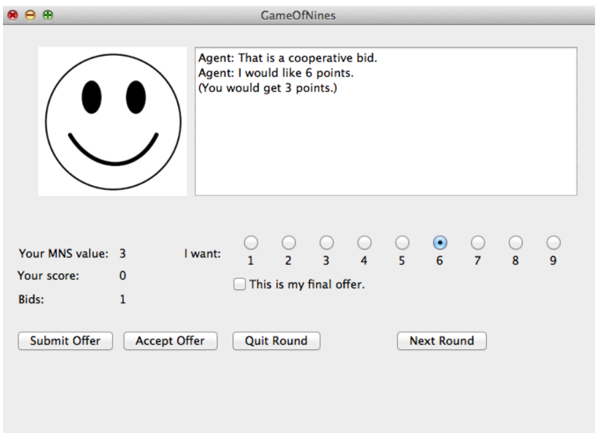

Figure 1: Used interface to play against the metacognitive agent.

2.2.2 Feedback

The agent could provide feedback to the human player about the strategy it was currently using. This feedback consisted of two cues given through the interface. The first cue represented the current strategy of the agent. Whenever the model used its cooperative strategy, a smiley face was shown in the left window (Figure 2). If the model was being aggressive, an angry face was depicted (Figure 2). This form of feedback was given to encourage subjects to be aware of the opponent's current strategy. The second cue was about the player's offer itself. The agent matched the player's action with its instances in declarative memory to determine whether the bid could be labelled as aggressive, cooperative or neutral. For every offer the player made, the agent notified the subject by saying "That is a cooperative / aggressive / neutral bid". This was again provided to give the player more insight into the agent's strategy.

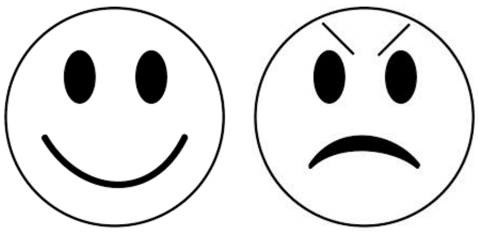

Figure 2: Faces used to show the metacognitive agent's current strategy. The smiley face (left) was shown when the agent used its cooperative strategy. The angry face (right) for the aggressive strategy.

2.2.3 No feedback

The agent could also play against a human player without providing feedback. In this case, no face was shown in the left window, which was left blank instead. The agent also did not notify the subject abouts its labeling of the player's offers.

2.3 Experimental structure

43 people participated in this experiment (13 male). Each were given €10 for their participation. Most participants were international students at The University of Groningen.

Before the experiment started, instruction sheets were handed to the participants who were then given sufficient time to read it through. This instruction sheet consisted of a detailed explanation of the rules of Game of Nines, an explanation of the interface used to play against the agent and two examples of a possible negotiation scenario in the game (Appendix A). Three high scores obtained in the pilot experiment were also included for motivational purposes. After each participant was finished reading the instructions, the supervisor read the instructions out loud and the participants were asked to read along. This was done to reduce the chance of a participant not understanding the rules. Each subject was allowed to ask questions at any time.

Once the instructions were clear, the participants moved to their own separate room where they would play against the metacognitive agent. At most three people performed the experiment simultaneously. The model was run on a Macbook Pro running OS X 10.9 Mavericks. The interface was shown in window mode on a 24 inch BenQ XL2420T monitor (16:9, 1080p). The participants used a regular mouse to control the interface.

All participants played against the metacognitive agent for two periods. Each period contained 3 blocks, consisting of 12 rounds each. The subjects were allowed a 5 minute break between periods, but were also free to start straight away. The following pairs of MNS values were used in each block: (1,1) (2,2) (3,3) (4,4) (1,3) (3,1) (1,5) (5,1) (2,6) (6,2) (4,5) (5,4). These pairs were used in Kelley's second experiment (Kelley et al., 1967). The order of these pairs were randomized for each block. However, the sum of the MNS values was the same for each participant. Thus no player had an advantage over any other.

In order to study the influence of feedback, the participants were split up into a control group (21 people) and a test group (22 people) without their knowledge. The test group played against the model with feedback in the first period, but received no feedback in the second period. The control group played against the agent without any additional feedback in both periods, meaning that the conditions of the second period were the same for all participants.

The purpose of this experimental design was that the subjects in the test group were stimulated to distinguish the two strategies the agent was using in the first period, by implicitly showing it. No feedback in the second period was provided in order to study whether these participants would use the information they may had learned in the first period in the (non-feedback) second period as well.

3 Results

Out of the 43 participants, the data of two subjects were not used. This was due to the fact these subjects had a negative score, which indicated that the rules were not clear to them. This resulted in a dataset containing 21 subjects in the test group and 20 in the control group.

Figure 3 shows the mean human score of the test group and control group for both periods. Figure 4 is a more detailed graph in which the human score is plotted per block.

Figure 3 shows that both groups had a higher score in the second period compared to the first. Fitting a linear mixed-effects model shows that the period had a significant effect on the human score (Table 1). The test group had an average score of 33.4 points in the first period and 39.9 points in the second, compared to 33.9 points and 37.7 points for the control group respectively. This suggest that human players on average did seem to improve over time when playing against the metacognitive agent. Furthermore, 76.2% of the participants in the test group had a higher score in the second period than the first. This was true for 55% of the control group.

The results of Table 1 also show that the interaction between period and whether the participant was in the test or control group is not significant.

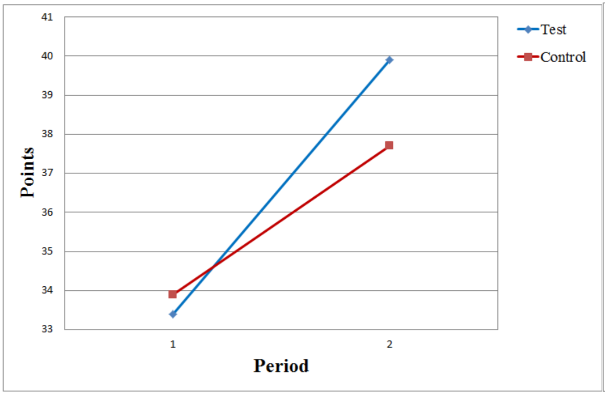

Figure 3: Average human score over period plotted for test and control group

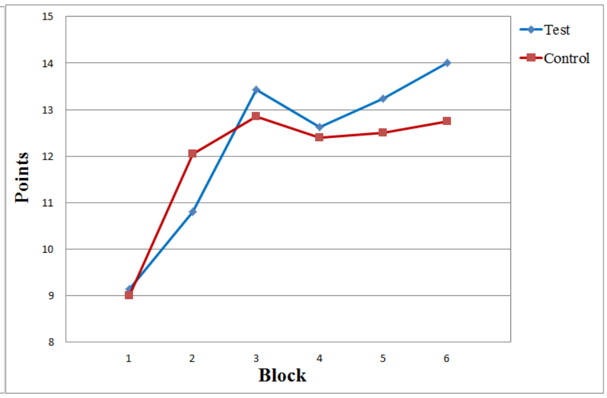

Figure 4: Average human score over block plotted for test and control group

Variable	Estimate	std.error	DF	T-value	P-value
(Intercept)	33.9	1.889	55.9	17.950	< 0.001
Feedback	-0.519	2.639	55.9	-0.197	0.8448
Period	3.750	1.627	39.0	2.305	0.0266
Feedback : Period	2.726	2.273	39.0	1.199	0.237

Table 1: Output of fitting a linear mixed-effects model for average human score per period (Figure 3)

This suggest that providing feedback to the participants did not have a significant effect on the score across all participants.

Figure 4 shows that both groups improve rapidly from block 1 to block 3 (period 1). A small decline in score can be seen at block 4, followed up by a slight improvement in block 5 and 6 (period 2). Overall, the participants from both groups appear to improve more in the first period compared to the second.

No other significant findings occurred when analyzing the control and test group. Thus the decision was made to split up the subjects from the test group and control group into medians based on their score in the first period. This resulted in two medians with 11 participants in the test group and 10 in the control group for the worst performing median and 10 subjects each in the test and control group for the best performing median. Figure 5 shows the score per period for the best and the worst performing participants.

Figure 5 shows that the best performing participants of the test group improved from an average of 39.6 points to 46.7 points, whereas the subjects of the control group did not improve (40.1 to 39.6 points). Fitting a linear mixed-effects model shows that there is a significant three-way interaction between feedback, period and the median (Table 2). This suggests that feedback had an effect on the human score for only the best performing participants. For the worst performing subjects, all participants improved over time significantly, but feedback did not seem to have an influence on the human score.

Furthermore, 80% of the subjects in the best performing test group had a higher score than the agent in second period. This was true for 30% for the control group. For the lower median, 30% and 10% of the test and control group respectively beat the agent.

Figure 6 shows the proportion of the agent's use of the aggressive strategy over the time periods. The agent could either use its aggressive or cooperative strategy, thus the proportion of both strategies can be derived from this graph. The results for both groups of each median are plotted.

The figure shows that the rate of aggression was relatively stable over time for three groups, with the exception of the best performing subjects in the control group. For the latter, the rate of the

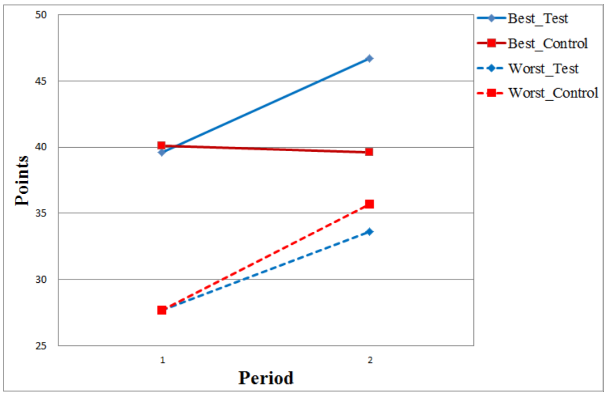

Figure 5: Average human score over period plotted for test and control group of both medians

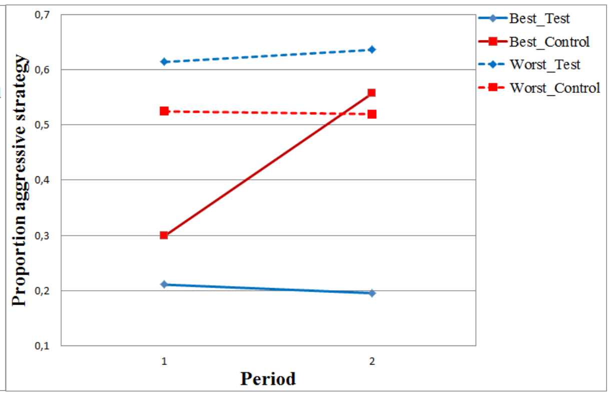

Figure 6: Proportion of agent's aggression against test and control group of both medians

Variable	Estimate	std. error	DF	T-value	P-value
(Intercept)	27.7	2.04	61.5	13.56	< 0.001
Feedback	0.027	2.82	61.52	0.010	0.992
Period	8.0	2.14	37.0	3.736	< 0.001
Median	12.4	2.88	61.52	4.293	< 0.001
Feedback : Period	-2.09	2.96	37.0	-0.707	0.484
Feedback : Median	-0.527	4.04	61.5	-0.131	0.897
Period : Median	-8.50	3.03	37.0	-2.807	0.00793
Feedback : Period : Median	9.69	4.23	37.0	2.289	0.0278

Table 2: Output of fitting a linear mixed-effects model for average human score per period for test and control group of both medians (Figure 5)

agent's aggression almost doubled over time period (29.9% to 55.7%). However, fitting a linear mixed-effects model shows that there are no significant interactions to be found, indicating that this rise of aggression is not significant compared to the aggression rates of the other groups (Table 3). Although, it might be possible that its significance is obscured because of the many other interactions tested in this statistical test.

Figure 6 shows that the agent was being the least aggressive when it played against the best performing subjects of the test group (21% to 19%), meaning it was cooperative about 80% of the time overall. Overall, it appears that the agent used a more aggressive strategy against the worst performing subjects compared to its more cooperative approach versus the best performing subjects.

Figure 7 shows the average first offer made per period for the best and worst performing participants. The first offer was defined as the first offer made by the human player in each of the 36 rounds per period. The figure shows that the worst performing subjects made higher opening offers overall than the best performing participants. An increase over time can also be seen. The control group of the upper median shows a similar pattern, though lower first offers. The test group of the same median showed a stable opening offer across periods and had the lowest opening bid of all groups in the second period.

Figure 8 shows the average offer made by the participants, plotted for the upper and lower median. The first offer was included when calculating this average. The figure shows that the best performing participants made lower offers on average compared to the worst performing participants. The lower median also made higher offers over time, whereas the upper median made more stable offers.

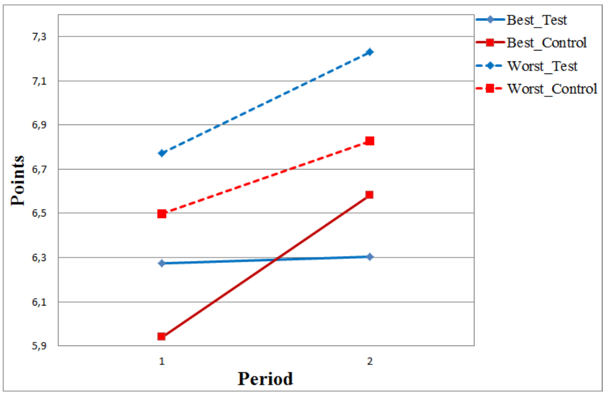

Figure 7: Average first offer over period plotted for test and control group of both medians

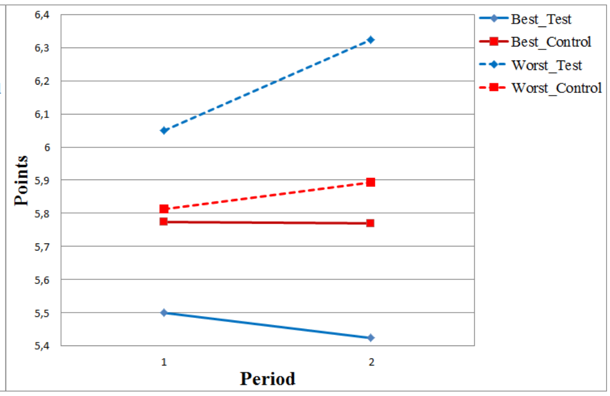

Figure 8: Average offer over period plotted for test and control group of both medians

Variable	Estimate	std. error	Z-value	P-value
(Intercept)	-0.432	0.694	-0.623	0.534
Feedback	1.483	1.017	1.458	0.145
Period	0.000187	0.947	0.000	1.000
Median	-1.044	1.0891	-0.959	0.338
Feedback : Period	-0.450	1.346	-0.335	0.738
Feedback: Median	-1.484	1.552	-0.956	0.339
Period : Median	1.479	1.424	1.039	0.299
Feedback : Period : Median	-1.868	2.174	-0.859	0.390

Table 3: Output of fitting a linear mixed-effects model for proportion of agent's aggression against test and control group of both medians (Figure 6)

4 Discussion

4.1 Implications of the results

The improvement in the score of the participants in both the test group and the control group (Figure 3) suggests that people did get better over time in Game of Nines by playing against the metacognitive agent.

The fact that no significant findings occurred for each logged variable between the control and test group suggests that feedback did not have a significant effect on the participant's performance in Game of Nines. However, the significant differences found between the best and worst performing subjects may indicate that only the best performing participants used the provided feedback to adapt their own strategy. This would explain the differences found between the test and control group of the upper median. It would also explain the similarity of the results of both groups in the lower median. If it is assumed that the subjects of this test group did not make use of the feedback they received in determining their own strategy, then it is unsurprising that their results are similar to the results of the control group, because their experimental conditions were then the same.

This difference can be seen clearly for the human scores for the groups in both medians (Figure 5). Figure 5 shows that the highest scoring participants of the test group improved significantly over time, whereas the control group did not improve further. This indicates that the feedback caused these participants to adapt their strategy in a way that improved their score. In contrast, the control group did not get to use this additional information and likely did not adapt their strategy to the agent's strategy and thus did not improve further.

However, the average score in the first period of the test and control group were the same. This means that the test group from the top median did not get a higher score than the control group when the feedback was actually provided, whereas the same participants did get a higher score in the second period (when the feedback was not provided). This may indicate that the feedback did not have an immediate influence on the participant's score, but that these participants still benefited from it in the second period. Yet, there is no clear evidence to support this causal attribution. Furthermore, the similar scores of the test and control group of the lower median suggests that these participants did not use the feedback in determining their strategy.

The other mentioned variables possibly show in what way the best performing participants used the feedback in changing their strategy. The results shown in Figure 6 show what strategy the agent was using against each group of participants. Additionally, it also gives information about the strategy the human player used. This is because the metacognitive agent aimed to imitate the strategy the human player used. This means that the aggression rate of the human participant was likely similar to the agent's.

From that information, it can be assumed that the highest scoring participants of the test group played the least aggressively and thus the most cooperatively. This suggests that the feedback caused these participants to play more cooperatively and that using this strategy results in a higher score when playing against the metacognitive agent. This was also suggested by Stevens (2015). This finding complies with the higher aggression rate of both groups of the lower median, which may have been a partial cause for the lower score of these groups.

However, the steep rise in aggression over time for the control group of the upper median does not fit with this explanation, because no significant decline in score is seen for these participants. This suggests that the strategy is not the only factor affecting the human score and that there are multiple ways to achieve a higher score.

The average first offers of all medians across periods (Figure 7) are in line with the observed strategy the participants from these medians used (Figure 6). A cooperative strategy is typically defined by moderate opening offers (Chertkoff and Conley, 1967). This complies with the low aggression rate and the low opening offers of the best performing test group. In contrast, an aggressive strategy is usually paired with high opening offers (Hüffmeier et al., 2014), which can be seen for the groups in the lower median. The increase of average first offers for the control group of the best performing median may be the reason the agent perceived their strategy as more aggressive over time period.

The results of the average offer of the four medians across periods (Figure 8) show similar trends compared to the average first offer. The best performing test group made the lowest offers on average, which is in line with its largely cooperative strategy. The average offer for the control group of the top median does not change over time period, suggesting that the rise in their opening offers was the main cause of the agent perceiving their strategy as more aggressive over time. Also, the higher offers that both groups of the lower median make are in line with their aggressive strategy.

Additionally, the test and control group saw a small decline in score when progressing from block 3 to block 4 (Figure 4). The break between these blocks may be a reason for this decline. Although most participants chose to start with the next period without taking a break, some time was still needed to start up the new session. The subjects may have needed some time at the start of block 4 to get into the game again. Furthermore, the test group shows a slightly larger decline on average compared to the control group. This may be due to the change from feedback to no feedback for these participants. However, it needs to be noted that both these findings are small and could be solely caused by error.

In conclusion, the results suggest that people did get better in negotiation by training against a metacognitive agent, but that only the best performing participants made use of the feedback and adapted their strategy to the agent's strategy. These participants largely used a cooperative strategy by making lower opening offers and lower offers in general. As a result, these subjects had the highest average score in the second period across all participants.

4.2 Critical notes

A number of aspects of this study can be revised. Firstly, all participants were students from the University of Groningen, meaning that selection bias may have occurred. Thus the obtained results may not be representative for the general population.

Secondly, the number of games the participants played may not have been sufficient to see an effect of the feedback for all participants. The learning rate in negotiation differs among people, meaning that some individuals could have benefited from playing more games (Cross, 1977). Thus it is possible that the subjects in the test group of the lower median might have also used the feedback to adapt their strategy if the experiment contained more games with feedback. Perhaps these participants had more trouble understanding the game mechanics and needed more time to develop a strategy, which left them with less resources to make use of the received feedback. This is supported by a study by Zohar and Peled (2008), which showed that low-achieving students needed a longer period of time in learning metastrategic knowledge to reach their top performance compared to high-achieving students.

Furthermore, the median split based on human score was done after data collection and all medians only contained 10 to 11 participants each. This undermines the statistical significance of these results and leaves more room for error.

4.3 Future research

This study suggests that people do improve at negotiation by training against a metacognitive agent and that people benefit from feedback on the agent's reasoning when the basics of the negotiation game are understood. For future research it would be interesting to study whether a metacognitive agent could be used for training in other tasks as well. Providing feedback on the agent's reasoning may also be useful for these other tasks. Further research should study whether metacognitive agents could be used as tools for improving people's skill. Training people in this way, as opposed to with conventional human training, may save a significant amount of effort.

Furthermore, for possible follow-up research, an experiment with more games and more participants is advised in order to study if these factors have an influence on the effects of feedback on human performance in negotiation. A possible experimental setup would be to let participants play against the metacognitive agent for 3 periods in which feedback is only given in the second period for the test group, meaning all participants have a full period to get to know the basics of the game. This way a greater number of people might make use of the feedback than in this study's experiment.

Another interesting study could be achieved by adding a feature to the game that enables participants and the metacognitive agent to make additional deals to each other. For example, the human player could propose a deal to the agent which says they will be more generous in their concessions in the next round, providing that the agent will be more generous in the current round in return. By adding such a deal-making feature, the negotiation game would become more closely related to real life negotiation scenarios.

References

Nancy J Adler. International dimensions of organizational behavior. *The International Executive*, 28(1):31–32, 1986.

David W Aha, Dennis Kibler, and Marc K Albert. Instance-based learning algorithms. *Machine learning*, 6(1):37–66, 1991.

John R Anderson, Daniel Bothell, Michael D Byrne, Scott Douglass, Christian Lebiere, and Yulin Qin. An integrated theory of the mind. *Psychological review*, 111(4):1036, 2004.

Otomar J Bartos. Modeling distributive and integrative negotiations. *The annals of the American academy of political and social science*, pages 48–60, 1995.

Jerome M Chertkoff and Melinda Conley. Opening offer and frequency of concession as bargaining strategies. *Journal of Personality and Social Psychology*, 7(2p1):181, 1967.

John G Cross. Negotiation as a learning process. *Journal of Conflict Resolution*, 21(4):581–606, 1977.

Joachim Hüffmeier, Philipp Alexander Freund, Alfred Zerres, Klaus Backhaus, and Guido Hertel. Being tough or being nice? a meta-analysis on the impact of hard-and softline strategies in distributive negotiations. *Journal of Management*, 40(3):866–892, 2014.

Harold H Kelley, Linda Linden Beckman, and Claude S Fischer. Negotiating the division of a reward under incomplete information. *Journal of Experimental Social Psychology*, 3(4):361–398, 1967.

Samuel S Komorita and James K Esser. Frequency of reciprocated concessions in bargaining. *Journal of Personality and Social Psychology*, 32(4): 699, 1975.

Raz Lin, Sarit Kraus, Yaniv Mazliah, et al. Training with automated agents improves people's behavior in negotiation and coordination tasks. *Decision Support Systems*, 60:1–9, 2014.

Janet Ed Metcalfe and Arthur P Shimamura. *Metacognition: Knowing about knowing.* 1994.

Dean G Pruitt and Peter J Carnevale. *Negotiation in social conflict.* Thomson Brooks/Cole Publishing Co, 1993.

Christopher A Stevens. Cognitive model of the game of nines. *Paper in preparation*, 2015.

Lawrence Susskind and Jason Corburn. Using simulations to teach negotiation: Pedagogical theory and practice. *Simulation und Planspiel in den Sozialwissenschaften. Eine Bestandsaufnahme der internationalen Diskussion, Münster*, pages 63–89, 2000.

Anat Zohar and Bracha Peled. The effects of explicit teaching of metastrategic knowledge on low-and high-achieving students. *Learning and instruction*, 18(4):337–353, 2008.

Appendix A

Instructions for Negotiation Experiment

Thank you for participating in this experiment. Today you will be playing a game in which you must negotiate with another player to divide up points between you. The goal of the game is to get as many points as possible for yourself in order to get a higher score.

The game you will be playing is called Game of Nines. On every round, you and your opponent will need to divide up 9 points between the two of you. Notice that, because there is an odd number of points, it is impossible to divide them up evenly. The more points you receive, the fewer your opponent receives and vice-versa. Each round starts by you (or your opponent) making an offer to the other. This offer consists of the number of points you want to receive (so 1 up to 9). Your opponent can then either accept the offer or do a counteroffer. Then it's your turn again and you can again accept the offer or make a counteroffer. The round continues until you or your opponent accepts an offer, or if one of you quits. In case of an agreement, you will each receive the amount of points you agreed on. If one of you quits however (thus no agreement has been made), both of you do not receive or lose any points.

Additionally, each round you and your opponent will receive a so called Minimum Necessary Share value (MNS value). The MNS value is the minimum amount of points you have to receive in order to add points to your score. For example, suppose your MNS value is 3. If you made an agreement in which you receive 5 points, then the amount of points added to your score will be 5 - 3 = 2 points. On the other hand, if you receive less points than your MNS value, you will lose points from your score. For example, receiving 2 points with a MNS value of 4 leads to losing 2 points of your score (2 - 4 = -2). Keep in mind, if you choose to quit, you and your opponent will get 0 points. Your opponent does not known your MNS value at any time. On a given round, your MNS value could be the same as your opponent's or it could be different. However, over the entire experiment, the sum of each player's MNS value is the same. So one player does not have an advantage over the other.

In this experiment, a computer model will serve as your opponent. You will go into one of the cubicles and use the interface pictured below on a computer to play against it.

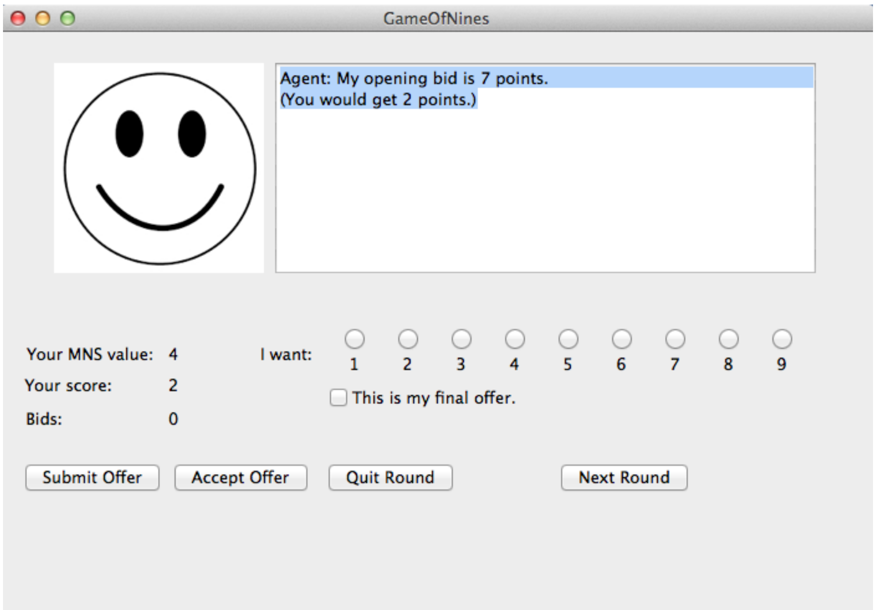

The computer model, referred to as agent, will communicate to you in the chat window. In the screenshot above it's making an offer saying it would like to receive 7 points. The line below that shows how many points you would receive if you were to accept its offer (9 - 7 = 2). Keep in mind you have to subtract the MNS value yourself to get the amount of points added to your score. You (and your opponent) have the option to make a final offer. If you choose to do so, your opponent has to accept or reject your offer and it cannot make a counteroffer.

There is no limit on the amount of offers you would like to make in a round. However, you will have 1 minute to make your move. If you don't submit an offer or quit in 1 minute, the round will end automatically and both you and your opponent won't receive any points.

You will play in 2 periods, containing 3 games of 12 rounds each. The model will then tell you to contact me ("Please contact supervisor") and you are allowed to have a 5 minute break. Then you will play the second period. The total duration of the experiment will be between 30 to 40 minutes. If you have any questions or if something isn't totally clear to you, please feel free to ask me.

Here are the top 3 high scores of earlier participants (total amount of points over one period)

1. 47 points
2. 44 points
3. 42 points

Try to get as many points as possible!

Good luck!