The Influences of Theory of Mind on Negotiation Performance in the Game of Nines by Use of a Meta-Cognitive Training Agent

Bachelor's Project Thesis

Isabelle Labat-Rochecouste, s2447932, belleroc@eircom.net,
Supervisors: Prof. Dr. N. Taatgen & C. Stevens PhD

Abstract: This study addresses the concept of theory of mind in terms of negotiation. Using a meta-cognitive training model our research looks at whether training on a negotiation agent that uses theory of mind will improve results in other negotiation tasks, in this case the game of nines. Participants in this study performed game of nines pre- and post-tests, and then trained with either a meta-cognitive training model with theory of mind or a training model without theory of mind. The difference in scores between pre- and post-tests for the game of nines were analysed. The results of the experiment are insignificant. This indicates that the presence of theory of mind in a training model does not seem to influence performance in these negotiation tasks.

Keywords: theory of mind; negotiation; training; game of nines;

1 Introduction

Negotiation has always been a core instrument in business and it continues to grow in importance both in professional and personal settings. Whether it's lawyers in talks for their clients, shop owners haggling with their suppliers or business men in meetings with their colleagues, negotiation is a vital part of modern society. This increasing valuation of negotiation has put a priority on teaching workers how to negotiate effectively and efficiently. For this reason, a negotiation training model is being developed to aid in the automation of negotiation training. Our research looks at the improvement and transfer of negotiation skills through the use of theory of mind.

Theory of mind can be quite simply defined as thinking about what other people are thinking. Theory of mind can be used on an emotional level which could be somewhat compared to the concept of empathy but it can also be used on a more tactical level when considering persuation and negotiation. We develop the capabilities of using theory of mind throughout infancy and it affects how we see the world and how we learn from it (Carlson, 2013). It has been shown through the 'false belief test' that between the ages of three and five most children begin to understand that other people view the world differently than they do and that everyone has different perspectives (Gopnik, 1988). Therefore, theory of mind is with us from quite an early age and it helps us learn how to act in many social situations from childhood through to our adult lives. However, it has been shown that we do not always use theory of mind in situations that could benefit from it (Hedden, 2002; Wright, 2010; Camerer, 2015). Previous research has also shown that negotiation training can improve people's negotiation skills, (ElShenawy, 2010; Movius, 2008) and also that thinking from the perspective of our opponent through use of theory of mind is a vital part of negotiation (de Weerd, 2015; de Weerd, 2014; de Weerd, 2013). As a result of this, the question follows as to how we can encourage people to use theory of mind to learn skills better and if negotiation itself benefits from theory of mind. Our research looks at whether theory of mind based negotiation training can show greater improvement in performance than training without the use of theory of mind.

Specifically, it looks at whether people who have trained with a meta-cognitive agent which encourages the use of theory of mind will show a greater improvement in a separate negotiation task, game of nines, than people who trained using a non-meta-cognitive agent which did not encourage the use of theory of mind.

To test this, two versions of the training model have

been created, a meta-cognitive version and a non-meta-cognitive version. The meta-cognitive version has been implemented to use theory of mind itself and to encourage its opponent to also make use of theory of mind. It assesses the tactics of its opponent based on factors such as how much information they provide about their own preferences in negotiation and how accommodating they are to the agent's requests. Based off of these factors the agent changes its own strategy to mirror that of its opponent's. For example, if the agent deems that its opponent is playing in an aggressive way (not sharing preferences, repeating proposals, dismissing the agent's suggestions) then the agent will play aggressively in return. On the other hand, if the agent labels its opponent's strategy as being co-operative (sharing their preferences, suggesting exchanges, considering the agent's suggestions) it will then adopt a co-operative strategy as well. The agent values recent moves higher than older moves as the impact of the opponent's behaviours on the agent's assumption decays over time.
The meta-cognitive agent encourages theory of mind by providing feedback to its opponent. This feedback consists of letting the opponent know which strategy the agent is using (aggressive, co-operative, neutral), letting the opponent know what the agent thinks their preferences are, and letting the opponent know what the agent thought of their last move i.e. whether they thought the move was aggressive or co-operative. When a round ends, the opponent is also informed of whether or not the previous deal was Pareto optimal. Pareto optimality means that the only way either of the players' scores could have been higher is if their opponent's score went down. A negotiation is considered unsuccessful if a deal other than the final deal was present which could have increased both players' scores. It is hoped that being exposed to this information will encourage the human player to use theory of mind and that they will adjust their strategy in accordance with the agent's strategy. The second version of the model is an agent which permanently stays in the co-operative strategy and that does not provide any feedback. This version is not intended to encourage the player to use theory of mind as they are not given any indication about the agent's thoughts.
The task we are testing for improvement in is the game of nines. The game of nines is another negotiation task however it is one-dimensional as opposed to the four-dimensionality of the training task. The game of nines agent is also meta-cognitive and players are informed about the current strategy of the agent by means of a happy, neutral or angry face, however no other feedback is provided. It is hoped that the participants who trained using theory of mind will use theory of mind in this task as well.
Our hypothesis is that participants who trained on the meta-cognitive model will improve significantly more in the game of nines from pre- to post-tests than the participants who trained on the non-meta-cognitive agent.
This experiment is done in collaboration with Marten de Jager who tested for an effect on human to human negotiation results. While both tasks will be mentioned in the following section, only the results of the game of nines task will be analysed and discussed in this paper. For the results of the human to human negotiation experiment it is advised that you consult his paper.

2 Methods

2.1 Participants

We ran our experiment on twenty-seven participants, four of which were from the pilot experiment. We used a group of volunteer participants and they were randomly assigned to either the control or experimental group. Thirteen of the participants were male. The participants' ages ranged between twenty and thirty. They had varying occupations however the majority of them were students in the University of Groningen. We ran experiments with two participants at a time as they were required to interact with each other during sections of the experiment. The participants were given monetary compensation of sixteen euro unless their partner did not show up and they chose to reschedule in which case they were given an additional five euro.

2.2 Materials

The experiment was split into five sections: a human to human negotiation pre-test, a game of nines pre-test, training, a game of nines post-test, and a human to human negotiation post-test.
In the game of nines, the interface of which can

be seen in figure 2.1, participants had to negotiate with an agent on how to divide nine points between them. The participant had to select how many of the nine points they wished to take and the model could either accept the offer, make a new offer, or quit the round. The model would quit the round if it decided that the participant was being unfair. The round ended when an agreement had been made or if either the model or the participant quit the round. When either the model or the participant quit the round neither of them received any points. In each round both the participant and the model had their own Minimum Necessary Share (MNS) value which is the minimum amount of points they had to attain in that round in order to avoid losing points. As an example, if the participant agreed with the agent that they would take five out of the nine points and their MNS value was two, then the participants score would increase by three. Both the participant's and the agent's MNS values changed after each round. In order to calculate how fair the participant was being the model averages its own MNS values over the course of the rounds. It knows that the participant's MNS values will average out to be approximately the same as its own. It is then able to calculate the participant's minimum goal for the round. It decided that the participant was playing aggressively if their value for that round was very high compared to their calculated minimum goal. The final scores of the block were the cumulative scores from each round in the block. The aim of the game of nines was for the player to get as many points as possible.

The training section of the experiment then began. In the training portion of the experiment, participants negotiated with an agent about their opinions on four different topics regarding smoking policies using a GUI on the computer. Policies that they were in favour of were highlighted in blue and policies that they were not in favour of were highlighted in red. The strength of the colour indicated how strong their like or dislike of the policy is. They did not know the agent's preferences and the agent did not know their preferences. The participant then negotiated with the agent about which policies to put into place until an agreement had been made on all four categories. With the experimental group, the agent was meta-cognitive meaning its play style corresponded with the participants play style, the interface is shown in figure 2.2. If the participant was being aggressive or co-operative, the agent would be aggressive or co-operative respectively. The agent's current approach was shown through a happy, angry or neutral face for co-operative, aggressive and neutral strategies. The experimental group was also given feedback about whether the deal in the previous round was Pareto optimal or not and they were also given information about the agent's understanding of the players preferences. In the control group the agent always played co-operatively and the participants were not given any feedback or information, which can be seen in figure 2.3.

For more detailed information about the inner workings of either the training program or the game of nines it is suggested to view the negotiation training model deliverables or the game of nines submission paper respectively (Stevens; de Weerd, 2016; Taatgen, 2015).

The participants also took part in a human to human negotiation task in which they had to carry out the same task as in training but with another participant instead of with a model. Their preferences were provided to them on paper and the conversations were be recorded using Audacity(R) (Team., 1999-2014) with the final agreement of each trial being marked on the preference forms by the pairs.

2.3 Procedure and Design

Our experiment began with the participants signing informed consent forms, filling in the payment forms, (See Appendix A) and being briefed on the task they were required to carry out. The pair of participants then took part in the human to human negotiation pre-test. There were six trials in the pre-test, with the first trial having a time limit of ten minutes and all remaining trials having a time limit of five minutes.

The participants were then brought to computers in their own individual cubicles so that they could complete the game of nines pre-test. The participants completed two blocks of the game of nines pre-test which took approximately ten minutes in total to complete.

Training then began on the negotiation model. The training segment lasted for nine trials and took on average thirty-five minutes to complete. The participants then completed the post-tests the game of nines and then finished the experiment with the

post-test for the human to human negotiation in the same pairs as before.
The total length of the experiment was approximately an hour and fourty-five minutes depending on the participants. Only the game of nines results and the training data was analysed.

2.4 Analysis

Our independent variables were which group the participants were in and the time variable. The group variable had two levels, control or experimental and the time variable also had two levels, pre- and post-test. The group variable was between subjects and the time variable was within subjects. The dependent variable was the scores of the game of nines trials. The scores were calculated and kept track of by the model.
Our data was analysed in R using the software RStudio (RStudio Team, 2015) and we tested for score differences in the pre- and post-test game of nines trials between the control and experimental group. We used a 2x2 Anova for analysis.

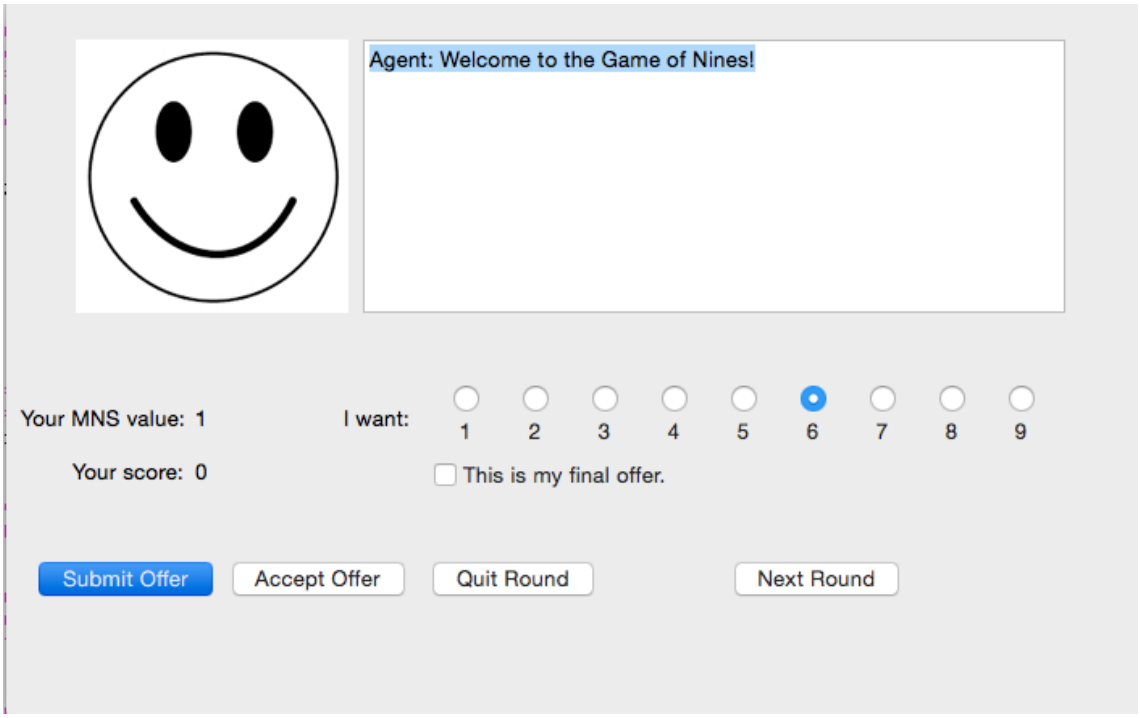

Figure 2.1: Game of Nines Interface

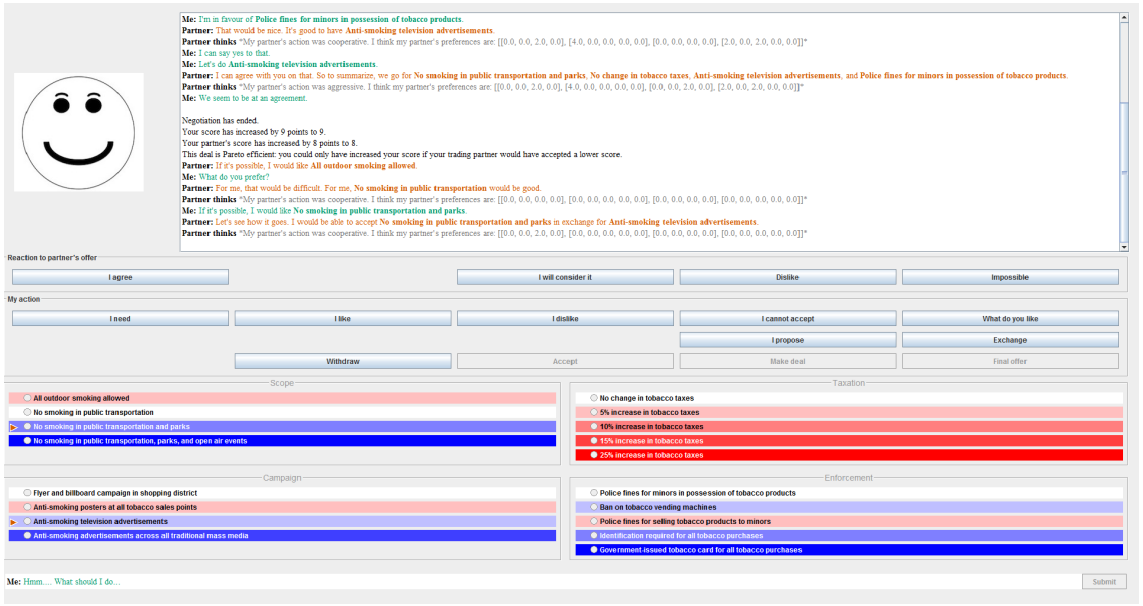

Figure 2.2: Meta-Cognitive Agent Interface

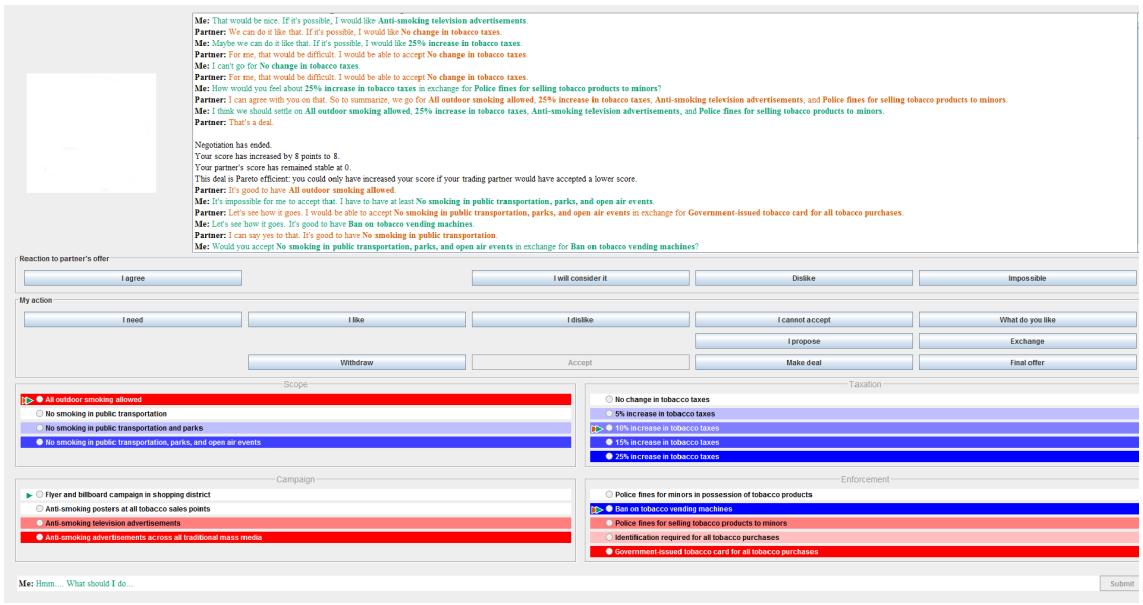

Figure 2.3: Cooperative Agent Interface

3 Results

A brief exploratory data analysis of Game of Nines scores between pre- and post-tests suggests that there is likely not a significantly greater improvement of scores in the experimental group compared to the control group. As can be seen in figure 3.1, the lines representing the control and experimental groups are almost parallel with the experimental line just being ever so slightly steeper. We found no statistical significance in our ANOVA, $F(1, 23) = 0.011$, $p = 0.9182$. The results of two participants had to be discarded. One of them had to be removed because they only completed one round of the game of nines post-test and the other one had to be discarded because they had previously completed an experiment which involoved the game of nines and we thought this might affect their results in some way. It should be noted that the mean scores of the control group are higher overall than the experimental group in both pre- and post-test.

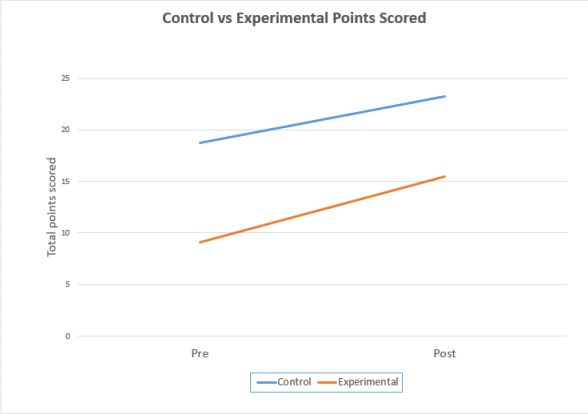

Figure 3.1: Interaction Plot of Game of Nines Scores Between Groups

4 Discussion

Our hypothesis was that participants who trained using the meta-cognitive agent which encourages the use of theory of mind would show a significantly greater improvement in Game of Nines scores between pre- and post-test than the participants who trained using the purely co-operative agent which did not encourage the use of theory of mind. While the interaction plot showed an ever so slightly steeper slope for the experimental group, our statistical analysis showed that there was no significantly greater improvement in the scores of the experimental group between pre- and post-test compared to that of the control group. Therefore, we must reject our hypothesis. This means using theory of mind did not lead to a greater improvement in negotiation performance scores in this experiment.

There are many possible reasons as to why there was no significant difference in improvement between the two groups. The first and most likely reason is that the participants did not train long enough with the agent to see an improvement in negotiation skills. The experiment generally took approximately an hour and fourty-five minutes to complete with training comprising thirty minutes of the total time. It is possible that this was not enough time for participants to get to grips with the training program and for them to benefit from it. The inclusion of two pre- and post-tests resulted in participants having very limited time to actually train with the model. If this experiment were to be repeated, I would suggest splitting it into two experiments. While this would take more time and resources, it would also allow for a far longer training period and could therefore show better results.

Another potential reason for the lack of greater increased scores in the experimental group is that they possibly did not actually make use of the theory of mind element of the training model. It was not explained in great detail to participants what the mechanisms of the meta-cognitive agent were in regards to the meaning behind the faces and the feedback so it's entirely possible that they did not pay attention to the written and visual feedback that the agent was giving to the participants. If this were to happen and the participants did not attempt to change their strategies in order to get better responses from the agent then it's likely that they did not interact with the theory of mind element of the training program. This would result in their scores not showing the effects of theory of mind when used by the participant but rather purely the effects of a meta-cognitive model.

A possibility in regards to the Game of Nines scores is that participants did not fully understand the task they were required to carry out. Many participants received negative scores in the Game of Nines which indicates that this possibility might

be true. If participants really didn't fully understand the task then the chance arises that they did not approach the Game of Nines from a strategic, negotiation based manner but rather in a more random, guessing game type approach. If this were the case it would explain why any possible negotiation skills learned during training did not transfer to the Game of Nines. It should be mentioned however, that when the negative scores were removed from the data pool the results were actually less significant than when the negative scores were included. This shows that the negative scores were not responsible for the insignificance.

As mentioned in the results section, the mean scores of the control group were higher overall than the experimental group. This could tell us that the control group were either better at negotiation from the beginning or that they pick up new tasks quickly. Therefore, it is possible that the control group simply took to the tasks they were asked to do better than the experimental group and as a result improved as much as the experimental group even without theory of mind based training. To counteract this, I would recommend observing the mean of the two groups during the experimental process and ensuring they stay approximately balanced.

Something else to consider is that perhaps the participants simply weren't ready to make use of theory of mind. If they found the training to be quite challenging as a task on its own then there's the chance that they did not have the capacity to use theory of mind on top of grasping the concept of the task they were required to do. This could be combatted by allowing participants to practice with a non-meta-cognitive agent for a couple of rounds before the actual training so they can become accustomed to the training task and layout.

A final potential reason is that theory of mind simply does not improve negotiation skills. While this explanation would not be ideal for the scope of the training model, it is still something that should be considered. Seeing as the control group benefitted the same amount from training without theory of mind and on a non-meta-cognitive model it could be said that theory of mind is not needed for improvement.

Future research could set out to amend the mentioned issues first and if the results show more significance then the experiment could be scaled up to include more participants. It would also be interesting to compare a computerised form of negotiation training with a classic human based training program to see if there are any differences in performance increase between the two.

In conclusion, it seems that autonomising negotiation training is a viable and potentially successful option which should be pursued further in the future. However, it is still unclear as to the importance of the inclusion of theory of mind into these training programs. If it is found that theory of mind continually does not cause a greater improvement in negotiation tasks, it would still be worthwhile to invest time into a training program that does not encourage theory of mind as it is indicated in our results that it still improves negotiation skills. On the other hand, if it is shown in future research that theory of mind does in fact result in a greater improvement then including theory of mind in negotiation training of all kinds should be encouraged.

Research should also be done on what other skills theory of mind can help train. As the world is becoming more and more autonomous in nature it is important to explore where it is useful to include technology in our lives and where it is not. Computerised, cognitive training programs can vastly improve the lives of many people as they could be more cost effective than a human based teaching program and therefore accessible to more individuals and can be included into people's schedules easier. They also have the potential to be personalised to the learning style of the individual, as high levels of education become more expected and more people are seeking out education it has become apparent that a standardised teaching system does not work for everyone. Computerised teaching and training could provide opportunities to people who had previously written off education such as those with autism, learning disabilities or who simply find it difficult to learn through classic teaching methods. In the future it is likely that many people will study through purely technological means so it is important that we put resources into finding the best and most effective forms of training for each field and specialisation. This is why it is important that we explore how theory of mind could improve education systems not just in negotiation but in other aspects of life as well.

References

Ho T. Chong J. Camerer, C. F. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898, 2015.

Koenig M. A. Harms M. B. Carlson, S. M. Theory of mind. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4(4):391–402, 2013.

Broers E. Verbrugge R. de Weerd, H. Savvy software agents can encourage the use of second-order theory of mind by negotiators. *In Proceedings of the 37th Annual Conference of the Cognitive Science Society.*, pages 542–547, 2015.

Stevens C. Taatgen N. de Weerd, H. Agent model of multi-issue negotiation. *The METALOGUE Consortium*, 2016.

Verbrugge R. Verheij B. de Weerd, H. Higher-order theory of mind in negotiations under incomplete information. *In PRIMA 2013: Principles and Practice of Multi-Agent Systems*, pages 101–116, 2013.

Verbrugge R. Verheij B. de Weerd, H. Theory of mind in the mod game: An agent-based model of strategic reasoning. *In Proceedings of the European Conference on Social Intelligence (ECSI-2014) CEUR Workshop Proceedings*, pages 128–136, 2014.

Eman ElShenawy. Does negotiation training improve negotiators' performance? *Journal of European Industrial Training*, 34(3):192–210, 2010.

Astington J. W. Gopnik, A. Childrens understanding of representational change and its relation to the understanding of false belief and the appearancereality distinction. *Child Development*, 59: 26–37, 1988.

Zhang J. Hedden, T. What do you think i think you think?: Strategic reasoning in matrix games. *Cognition*, 85:1–36, 2002.

H. Movius. The effectiveness of negotiation training. *Negotiation Journal*, 24(4):509–531, 2008.

RStudio Team. *RStudio: Integrated Development Environment for R.* RStudio, Inc., Boston, MA, 2015. URL `http://www.rstudio.com/`.

Daamen J. Gaudrain E. Renkema T. Top J. Cnossen F. Taatgen N. Stevens, C. Using a metacognitive model to improve negotiators theory of mind. *The University of Groningen.*

Stevens C. de Weerd H. Cnossen F. Taatgen, N. Cognitive model of a multi-issue negotiation task. *The METALOGUE Consortium*, 2015.

Audacity Team. Audacity(R) software is copyright (c), The name Audacity(R) is a registered trademark of Dominic Mazzoni., 1999-2014. URL `http://audacity.sourceforge.net/`.

Leyton-Brown K. Wright, J. R. Beyond equilibrium: Predicting human behavior in normal-form games. *In Proceedings of the Twenty-Fourth Conference on Artificial Intelligence.*, pages 901–907, 2010.

A Appendices

Form of Consent

You will participate in an experiment about negotiation.
In this study, you will play a negotiation game with one other participant.
The data will be stored and analyzed anonymously, to ensure your privacy.
Participation in this study is voluntary. You may quit at any time.
You will receive 16 euros for your participation.
The experiment will last roughly 2 hours.

Thank you for your participation
Marten de Jager
Isabelle Labat-Rochecouste

I agree to participate in the experiment on negotiation.

Name & Signature of Participant Date

Game of Nines Instructions

The game of nines is a game in which you must divide nine points between you and computer player with the aim being to get the most points.

In each round you will be given an MNS value on the left hand side of your screen. This value is the minimum amount of points you must get that round to avoid losing points. For instance, if your MNS is 2 and you agree on 3 points, your profit for the round is 1 and 1 point will be added to your score.

You can select how many points you wish to keep for yourself and offer this deal to the computer player. The computer player can then either accept your offer or make their own offer. If you cannot reach an agreement with the computer player, then neither of you get points.

If you take too long to respond the round will end and a new one will begin, if this is the case or if you have any other problems please inform one of the experiment coordinators.

Uitbetaling aan (buitenlandse) proefpersonen (FIOD)

Uitbetaald aan / Payed to:

Naam / Surname: _____

Voornaam / Firstname: _____

Adres / Address: _____

Postcode / Zipp code: _____

Woonplaats / City: _____

Fiscaal nummer: _____

Geb. Datum / Date of Birth: _____

Banknummer / Bankaccountnumber: _____

Bedrag Honorarium / Fee: _____

Bedrag Onkosten / Expenses: _____

Voor betaling aan buitenlandse proefpersonen aub ook de gegevens hieronder in vullen:

Name Bank:	
Bank Address:	
Zippcode:	
City:	
Country:	
BIC (Swift):	
IBAN code:	
Routing/Aba/BLZ:	
Currency:	

Voor de volgende verleende diensten: / for the following services:

_____

_____

Projectcode: _____

Groningen, _____

Handtekening voorzitter
basiseenheid/projectbeheerder,

_____

In te vullen door FBA

Crediteurnr	**Fact. datum**	**Fact. nr.**	**Bedrag**	**BTW**
				GBO

Projectcode	**Structuurdeel**	**Kostencode**	**Bedrag**	**Omschrijving**

Postbus 407
9700 AK Groningen