

GAME OF NINES: THE EFFECT OF ADDING MORE REALISM TO THE HUMAN-AGENT BARGAINING PROCESS

Maikel Grobbe, s2041278, h.h.grobbe@student.rug.nl

Abstract: Game of Nines is a game where 2 people bargain with each other to divide 9 points, where they both have a certain minimum necessary value(MNS) they have to achieve. This game is used, since it was first introduced by Kelly et al. (1967), to get a better understanding of the bargaining process. In the study performed by Daamen (2014) a start was made to make a simulation of the human-agent bargaining process. In this study we set out to add a layer of reality to this simulation by adding the ability to call the agent on their lies, asking for a favor and not having to specify your MNS as well as adding another agent which does not change behavior based on the opponent. Results show that these additions make for a simulation where people are more prone to change tactics based on their opponent. They do this by changing the amount they ask for a favor, the amount they use a final offer and the amount they lie about their MNS value.

1. Introduction

Bargaining is something that has people interested for several decades. What tactics people use and how they use them to get the best result is something that only the best salespeople truly understand, or at least master. Because our society keeps improving itself with technological advances, the question arises whether or not we could find something that truly gives us full insight in the way we humans bargain with each other and, furthermore, with agents. We are interested in whether or not we humans can distinguish between different strategies and react upon these in a good manner.

Since the 1960's a lot of research has been done on the topic of bargaining, with Kelly, Beckman and Fischer (1967) introducing the Game of Nines, the bargaining game that we use for this experiment, and other researchers like Liebert, Smith, Hill and Keiffer (1968) and Schoeninger and Wood (1969) adding onto this game with their different views on the subject. Recently, Mascarenhas, Marques, Campos and Paiva (2013) constructed a model of social dynamics to explain human behavior more extensively by being able to explain findings in a children's version of Game of Nines. Stevens (unpublished) did a person versus person study of Game of Nines, extensively marking the conversations between test subjects. This was then used in Daamen (2013) his study to try and simulate the human-agent bargaining process involving the Game of Nines game.

The Game of Nines game is a game played by 2 persons. They both get given a minimum necessary share (MNS) between 1 and 6 and then have to divide 9 points between them. They do this by bargaining, making promises, lying about information and leaving out information. Every point they get above their MNS value is the score for that round once an agreement is reached. Once an agreement is reached, either of the players calls for a stop or the time limit of 60 seconds is reached, a trial ended. For example, if player A has a MNS value of 3 and player B has a MNS value of 2, you might assume that the agreement they will reach will be at the share of 5 for player A and a share of 4 for player B, giving them both a score of 2 for that round. By lying about his MNS value, or asking for a favor, player B might convince player A that the division should be 4 for A and 5 for B, giving him a big advantage that round. This might however, because of players playing multiple games and having to build up trust over the course of the games, lead player A to become suspicious of B and netting him less points than in the scenario where he would have played fair. Exactly these choices that people make when adapting different strategies, possibly against different opponents, is what we are interested in.

In the experiment performed by Daamen (2013), he already searched for this. In this study subjects played against a fair and an unfair agent. Subjects had to say something about their MNS value, what they wanted and whether or not their current offer was their final one. In his experiment he found a significant change in the

use of the final offer depending on which agent a subject faced. We felt that in the end the simulation created by him was not extensive enough because of limited options for the players, which is why we are adding on the research already performed by him. We do this by looking at the experiment performed by Stevens (unpublished) and see what we are missing in the experiment performed by Daamen (2013). Analysis revealed that a big part missing in the experiment was the opportunity to ask the opponent for a favor, that is making promises about future hands. Next to this we also found that people sometimes called the opponent on a lie and not always said something about their MNS value, something that was mandatory in the Daamen (2013) experiment. By adding these functions and creating a third agent which does not look at what the opponent is doing, something we also believe to have found in the experiment by Stevens (unpublished), we hope to create an experiment that lives up to the job of a more realistic simulation of the human bargaining process. What we hope to find out in this experiment is whether or not subjects change their behavior when they play against the 3 different agents, which we do by trying to create the simulation as described.

2. Method

2.1. Subjects

For this experiment 19 participants have been used, of which 15 male and 4 female. The mean age of the participants was 20.5, ranging from 18 up to 26 years old. The participants were recruited by posting on the website of the study association for artificial intelligence and computing science at the RUG, study association Cover (svcover.nl).

Participants were offered 10 euro for participating in the experiment which took about 1 hour to complete including instructions.

2.2. Instruction sheet

Upon entering the experiment booth, participants had to read the instruction sheet before starting with the experiment. The instruction sheet gave an overview of the game they were about to play, how many hands they were going to play, how many agents they

would be facing and how many times they would be facing each agent. They were told the goal of the game was to score as many points possible.

After that they were given a view of the interface, explaining what options they had during the game. It was explicitly mentioned which options were mandatory and which were optional. It was mentioned that they would play each agent 3 times and that agents would reset between every instance that you played them, meaning they did not remember what happened then. No details were given about what each agent could or could not do and how they would react to the options that were given to the participant, as to not give them more information than needed and influence them.

They were asked before starting the experiment if everything was clear and whether they had any more questions. If they didn't, they had the permission to start the experiment which was already set up. If they had further questions during the experiment they could ask these as well.

2.3. Simulation

The Game of Nines simulation as used during the experiment takes the simulation as written by Daamen (2013) and enhances on this. It is written for Python 3.4 and uses the tkinter toolkit for constructing the user interface (UI). During the simulation participants played against 3 different agents. These included Tom, an agent who wants to make fair deals, Ben, an agent that wants to get as many points for himself using lies and other tricks, and Jouke, an agent that does not look at the opponent at all and has a list of values he wants for certain MNS values. Further details about these agents will be given later in this paper.

Each agent had a face next to their output in order for the participant to better differentiate between the certain agents. These faces were randomized between all experiment booths in order for them to be different amongst test subjects, as to not have any effect on the experiment.

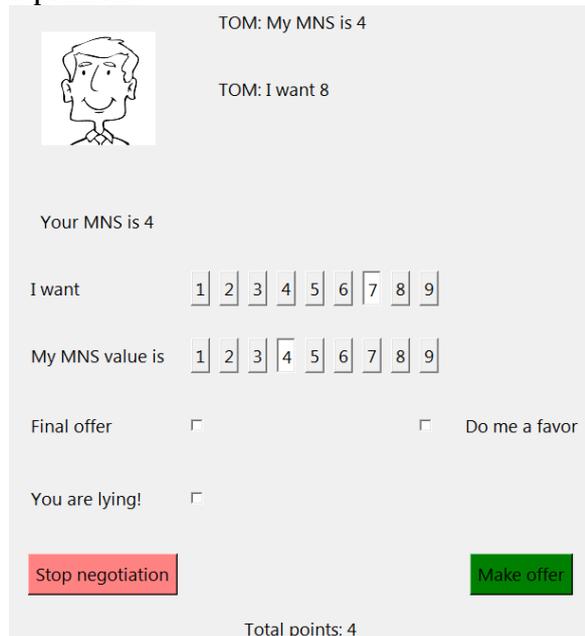
The participants were randomly divided into 6 groups with varying order of the 3 agents, meaning that every order possible was used during this experiment. The first order (Tom,

Ben, Jouke) was used 4 times, all the other orders were used 3 times.

Each participant played each agent 3 times, making a total of 9 blocks played per participant. Within each block 12 MNS value sets were played. The MNS values were played in a random order, meaning that every time you played an agent the order would be different. This meant that participants had to think ahead when choosing certain actions, deciding if it would also benefit them in the long run. The chosen values represent a distribution of low and high MNS values. the average of these MNS values is the same for both the agent and for the participant. The values that were used are the following, shown as (MNS value for participant, MNS value for agent); (2,2), (1,3), (3,1), (2,2), (3,3), (2,3), (3,2), (3,4), (4,3), (2,4), (4,2), (4,4).

The interface that participants communicated through during the experiment is shown in figure 1.

Figure 1: User interface for the Game of Nines experiment.



At the start of each trial the participant must say something about the division he wants for the 9 points of that trial. He does this by clicking on the amount of points he wants for himself. In the example shown in figure 1, the participant chooses to ask for 7 points, meaning he is willing to give the agent 2 points.

After the player has selected the number of points he wants, he can then select whether or

not he wants to say something about his own MNS value, that is shown beneath the picture of the agent. If he chooses to do so, the agent will also say something about his MNS value, as shown on the first line of figure 1. He does not have to do this.

Furthermore the participant can find 3 optional actions beneath that. He can say that the offer he is making is his final offer. If the agent accepts it, the current offer will result in an agreement. If the agent does not accept it, the current trial will result in a negotiation stop and give both players a score of 0 for that round. If a participant uses the 'do me a favor' button, the agent may or may not agree to grant a favor. Whether or not the agent grants a favor depends on which agent you are playing. The last option a participant has to influence the agent is to call the agent a liar. How the agent reacts also heavily depends on which agent you are playing. How each agent reacts to the different options a player has is described in the next section.

Negotiation about a given trial ends in one of the following 4 circumstances; An agreement is reached, and the score gained by each player is added to their total score. The player clicks on the 'Stop negotiation' and the negotiation is stopped, resulting in 0 points for both the participant and the agent. The agent or the participant rejects a final offer, resulting in not getting an agreement and getting awarded 0 points each. The trial time-out of 60 seconds is reached resulting in both parties getting 0 points.

2.4. Agents

Tom, the fair agent

The fair agent Tom tries to always get a good deal for both himself and for the opponent. He never lies and tries to be fair all the time. He does this by trying to come to a even division of the points that are to be gained. This means that for example if the player has an MNS value of 3 and the agent has an MNS value of 2, he will try to come to an agreement at 5 for the player and 4 for the agent, giving them both a score of 2 for that round.

This is what Tom always tries to do as long as he thinks that the player is not lying to him. Tom has 2 equations to determine whether or not the opponent is lying to him, by generating a

trust value. This trust value can be seen as a way for the agent to determine whether or not the opponent is lying and if he is lying how much he is lying. He does this by looking at the difference between MNS values of the opponent and itself or the first offers with which an opponent started compared to the offer it had in mind.

The first equation is based on a situation where the player has said something about his MNS value and the second equation is for when the player has not said anything about his MNS value. These equations don't always work though, seeing as at the start of each block, the agent has no information to compare between his own MNS or offer pattern and that of the opponent. This is why equation 1 and equation 2 are only used after a certain amount of turns. Equation 1, which uses the MNS value of the player, is only used once the player has said something about his MNS 5 times. Equation 2, which uses the first offers of a player in a block and compares that to the first offer that the agent made in that block to assess a trust value, is only used once the players have played more than 3 hands. In the case that neither equation is used, a trust value of 0 is asserted to the player, which means that a player is trusted completely.

If the player has given no information about his MNS value, the agent guesses a value in which range it will probably be. It does this by using equation 3. Equation 3 takes the maximum of 9 points and then subtracts his own MNS value from that, seeing as that is the maximum that the opponents player MNS can be. It then divides this by 2 leaving you with the average MNS value that the player will have given your own MNS value.

Once the trust value is calculated using equation 1 or 2, equation 4 is used to determine the adjusted MNS with which Tom then determines a fair offer. Since Tom tries to always strike a fair deal, he does not like people lying to him, making his agent type adjustment 1.

Equation 1: generating trust value given a MNS value.

$$\text{trust} = \frac{\text{diffMNS}^2}{\text{se(MNSlist)}}$$

Equation 2: generating trust value given no MNS value

$$\text{trust} = \frac{\text{diffFO}^2}{\text{se(FOlist)}}$$

Equation 3: guessing what the MNS value of the opponent might be

$$\text{gMNS} = \frac{(9 - \text{oMNS})}{2}$$

Equation 4: adjusting the MNS value based on trust value

$$\text{aMNS} = \text{cMNS} - \text{a} * \text{trust}$$

In these equations the following variables are used;

trust = the value calculated for trust.

diffMNS = the difference in MNS means between the agent and the player

MNSlist = a list with all the values for the MNS used by both agent and player

diffFO = the difference in first offer means between the agent and the player

FOlist = a list with all the values for the first offer used by both agent and player

gMNS = the guessed MNS of the player

oMNS = the MNS of the agent

aMNS = adjusted MNS value for the player

cMNS = claimed MNS value by the player

a = agent type

Tom reacts to a final offer in the following manner. If the offer is a fair offer, he will always accept. If the offer is one point above a fair offer he has a chance of 30% to accept the offer. If the offer is 2 points above a fair offer, he has a 10% chance to accept the offer. If the offer is 3 points or more above a fair offer, Tom will never accept and stop the negotiation. After giving 2 counteroffers, Tom has a 50% chance to ask for a final offer himself for every round after that one.

If Tom is asked for a favor, he will always agree to do this. He does this by lowering the fair offer that he has by one for himself, thus adding one to the fair offer of the player. The player will then thus score one point more if he comes to an agreement that round. If the player accepts this favor by coming to an agreement that round, Tom then seeks to get back this favor in the trials to follow. If in the rounds to follow a situation arises where the offer of the player is one above the fair offer as determined by Tom, he will ask the player for a favor. If then in the next round (meaning the next time the player makes an

offer) an agreement is reached, meaning the player has lowered his offer by 1, he sees this as honoring his question for the favor. If no agreement is reached the following turn, he will try once more to ask for a favor, that is, if the situation has stayed the same. After asking for a favor twice and not receiving it he will stop asking for any more favors. If the favor of Tom is returned, he will then be open for the player asking another favor. If the favor of Tom is not returned, he will not give a favor to the player anymore.

Since Tom does not lie and always tries to strike a fair deal with the player, he will not react to any accusations made by the player.

Ben – the unfair agent

The unfair agent Ben has no interest in playing fair and only tries to maximize his own score by trying to trick the player into believing he is giving him a good deal and playing very aggressively.

For determining a fair offer, Ben also uses equation 1, 2 and 3. In contrary to Tom, Ben uses a value of 0.5 for his agent type. He does this because he is lying himself, meaning that he is not as punishing against the opponent lying. To determine how much he wants to lie, he uses equation 5 to determine a new MNS value for himself, on which he then determines a 'fair' offer. If the opponent says something about their MNS value, he also states this adjusted MNS as his MNS value. In equation 5, cMNS is the claimed value of Ben.

Equation 5: determining an adjusted MNS value

$$cMNS = MNS + \frac{4}{MNS}$$

Ben also has a larger chance of accepting final offers. He bases this on how much he can gain from a final offer that is given to him. If the offer is 3 or more points higher than his MNS value, Ben will always accept. If the offer is 2 points higher than his MNS value he has a 90% chance of accepting. If the offer is 1 point above his MNS value he has a 70% chance of accepting this offer. If the offer is exactly his MNS value he will never agree.

Ben mainly follows the same principles as Tom when it comes to giving favors. In the

beginning he will wait for the player to ask for a favor just as Tom does. If asked for a favor he will also directly offer one by lowering the fair offer that was determined by one. Instead of the 2 turns that then Tom tries to get his favor back, Ben keeps trying this for 5 times, unless his favor is returned. Furthermore, once he has gotten his first favor back, he will try to get a second favor out of the player, with a maximum of 5 tries again. If he succeeds to get a second favor out of the player, he will again be open to giving a favor himself.

Ben does react to people calling him a liar, seeing as he is constantly lying. If someone calls him a liar, he will, for 5 offers, use the same principles as Tom when it comes to assessing trust and making a fair offer, meaning that he plays fair for those 5 offers that he does. If during these 5 offers he is once more called a liar, he will go back to lying for the amount of turns that he still had left to play fair. After those turns are over he will play fair again for 5 turns until these are depleted. To state it more directly; If the value of the turns he is supposed to play fair is between 0 and 6, he will play fair, otherwise he will play unfair.

Jouke – the stupid agent

The stupid agent Jouke is very different when compared to the other 2 agents that are used in this experiment. In contrary to the other agents he does not look at what the other player is asking, because he is just not interested. He is not per definition really stupid, he is just oblivious too everything that is happening around him and does not act on the information that is given to him.

To determine how many points he wants to get, he only looks at his own MNS value and bases what he asks upon this. He follows the table 1 for determining his offer. Furthermore, his claimed MNS is the amount of points that he wants in his offer minus 1, as also stated in table 1.

Table 1: Jouke's offer table.

Jouke's MNS	Jouke's offer (player, agent)	MNS value that Jouke tells his opponent
1	(4, 5)	4
2	(4, 5)	4
3	(3, 6)	5
4	(3, 6)	5
5	(2, 7)	6
6	(2, 7)	6
7	(1, 8)	7
8	(0, 9)	8

In contrary to the other agents, Jouke will accept an offer that is one below what he is asking, This means that if he wants 5, but is offered 4, he will accept without any hesitation.

When Jouke is told that something is the opponents final offer, he will behave exactly as Ben, meaning a 90% chance on 2 points above MNS, and 70% on 1 point above MNS. If he does not gain any points he will not agree to a final offer.

When he is asked for a favor, he will follow the same principles as Tom, meaning he will give it without any problems and then try to get it back for 2 turns after that.

Jouke does behave a bit differently when it comes to calling him a liar. He follows the same basic idea as Ben, but with much less impact. Instead of 5 offers, Jouke will lower his offer by 1 for only 2 offers. If he is called a liar when already having a lower MNS value, he only adds on 2 also, and plays fair between 0 and 3 turns, having already 'forgotten' what the player said relatively quickly.

2.5. Data collection and analysis method

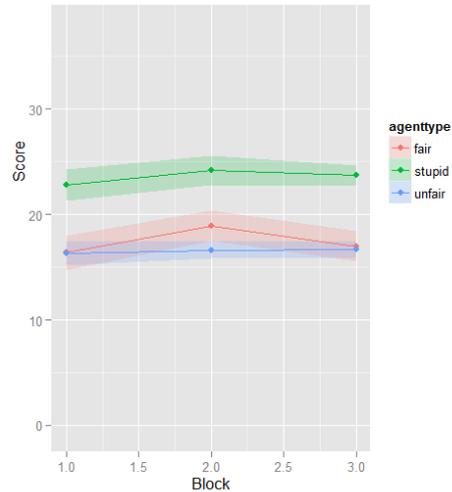
The data that was collected were the following variables for each and every turn of the game. Agent order, current agent, subject, age, sex, time, MNS of player, claimed MNS of player, offer for player, final offer use player, favor use player, lie use player, MNS of agent, claimed MNS of agent, offer for agent, final offer use agent, favor use agent, what action was done, person doing the action, number of total points, trial number, current fair offer that the agent proposes and the block number. In the programming phase an error was made, meaning that during the experiment the times

that the player called the fair agent Tom a liar were not correctly placed in the data logfile, meaning that a comparison between the different agents and the fair agent there is not available.

All the analyses found in the results section were done by using R. Furthermore the packages lme4, ggplot2 and lmerTest were used.

3. Results

3.1. Score

Figure 2: Score per block**Table 2: Linear mixed model fit of block and agent type on score.**

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	16.8070	1.7589	9.556	<2e-16	***
Block	0.3158	0.7564	0.418	0.6769	
Agenttypestupid	5.8421	2.3108	2.528	0.0125	*
Agenttypeunfair	-0.6316	2.3108	-0.273	0.7850	
Block:Agenttypestupid	0.1316	1.0697	0.123	0.9023	
Block:Agenttype unfair	-0.1316	1.0697	-0.123	0.9023	

What we test here is whether or not people improve their score over the time of the experiment. This is done by looking at a plot of the score over the blocks and further analyzing this with a linear mixed model that tries to explain the variance in the score value in terms of effect from the agent type and the block.

What we see is a significant effect from the stupid agent on the score both in the plot and the linear mixed model.

3.2. MNS difference

Figure 3: Difference between the real MNS value and the claimed MNS value of a player.

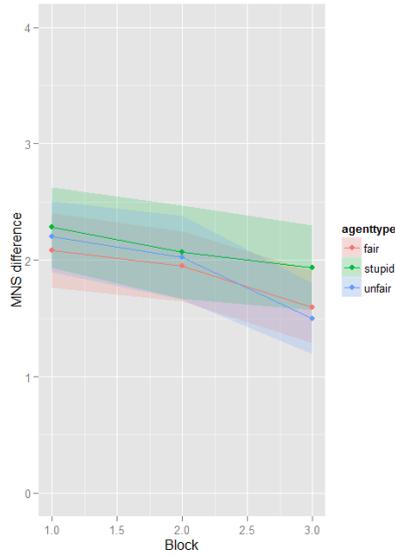
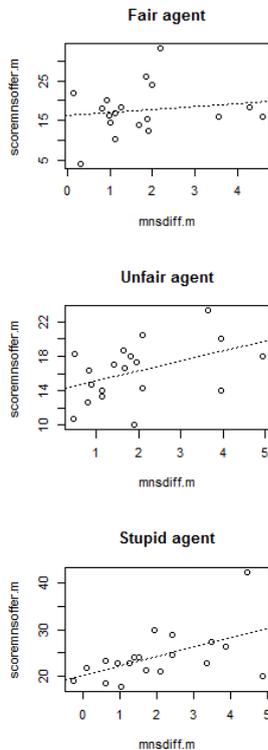


Figure 4: The correlation between score and lying about the MNS value



We look at whether or not people lie more about their MNS value over the time of the experiment. We do this by looking at the difference between MNS value given to the player and what they tell the agent their MNS value is and how this changes over the blocks. Furthermore we look at the effects of the different agents and the blocks on how much is lied about the MNS value with a linear mixed model. Next to that we have some correlations between score and lying about the MNS value.

The correlation of the difference in MNS and the score in the fair agent (p-value = 0.5532) and the unfair agent (p-value = 0.06041) are not significant. the correlation of the difference in MNS and the score is significant in the stupid agent (p-value = 0.02104)

Next to that we see that there is a significant effect from block on the amount people lie about their MNS value.

3.3. Final offer use

Figure 5: The use of final offer.

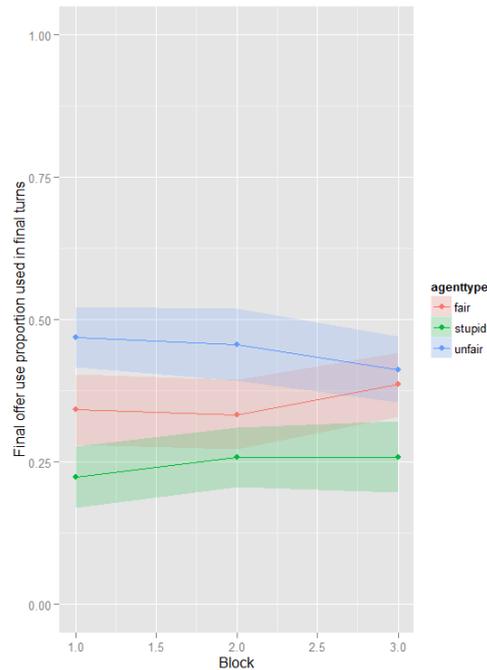
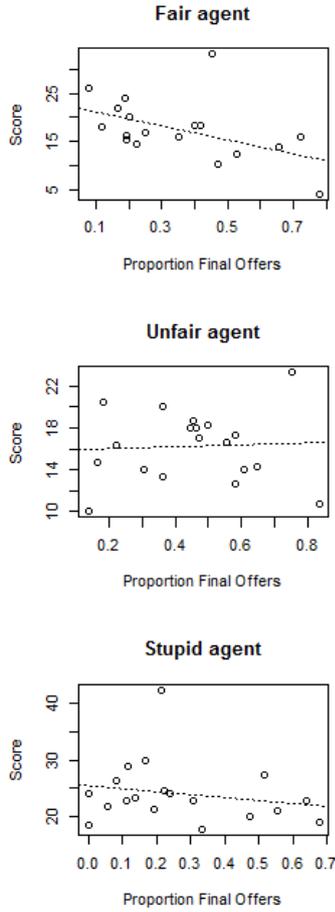


Table 4: Generalized linear mixed model fit of block and agent type on final offer use.

	Estimate	Std. Error	z value	Pr(> z)	
(Intercept)	-0.92555	0.312056	-2.966	0.00302	**
Block	0.113406	0.106143	1.068	0.28533	
Agenttypestupid	-0.59426	0.344843	-1.72	0.08484	.
Agenttypeunfair	0.956164	0.319225	2.995	0.00274	**
Block:Agenttypestupid	-0.00359	0.157922	-0.02	0.98183	
Block:Agenttypeunfair	-0.25005	0.147339	-1.69	0.08967	.

Figure 6: The correlation between score and final offer use



We are interested in whether or not people use the final offer option more when progressing in the experiment. We do this by plotting the amount of final offer use against the blocks and using a generalized linear mixed model to look at the variance in the final offer use caused by either block or agent type. Next to that we look at the correlation between score and final offer use in the different agents.

The correlations show a significant negative effect of score and final offer use for the fair agent (p-value = 0.0381) and no significant correlation for the unfair (p-value = 0.8379) and the stupid agent (p-value = 0.4097).

In the plot and the generalized linear mixed model we can see a significant effect from the unfair and the stupid agent on final offer use as well as significant interaction between block and the unfair agent.

3.4. Favor use

Figure 7: The use of asking for a favor.

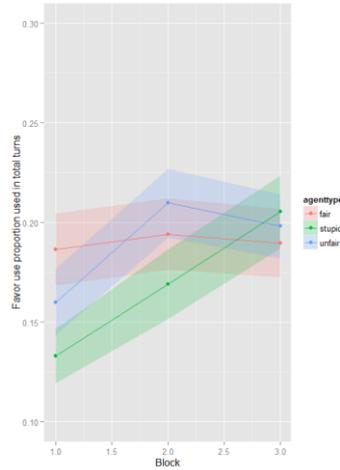
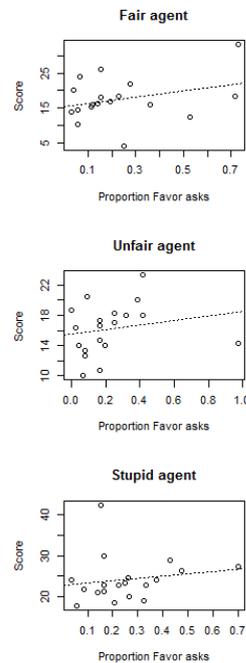


Table 5: Generalized linear mixed model fit of block and agent type on favor use.

	Estimate	Std. Error	z value	Pr(> z)	
(Intercept)	0.1869	4.04e-02	4.623	5.5e-05	***
Block	0.0015	9.84e-03	0.162	0.87133	
Agenttypestupid	-0.090	3.01e-02	-3.00	0.00274	**
Agenttypeunfair	-0.035	3.01e-02	-1.18	0.23525	
Block:Agenttypestupid	0.0346	1.39e-02	2.489	0.01289	*
Block:Agenttype unfair	0.0174	1.39e-02	1.256	0.20941	

Figure 8: The correlation between score and favor use



We do the same for favor use as we have done for the final use.

We see that there are no significant correlations between favor use and score in either the fair agent (p-value = 0.2172), unfair

agent (p-value = 0.4163) or the stupid agent (p-value = 0.5095).

We do see a significant effect from the stupid agent and the interaction between the stupid agent on the amount of favors used as can be seen in table 5.

3.5. Lie use

Figure 9: The use of calling opponent on a lie

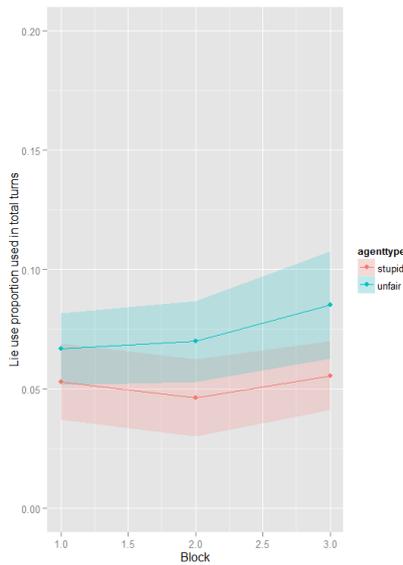
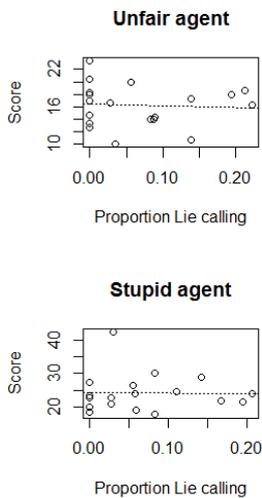


Table 6: Generalized linear mixed model fit of block and agent type on lie use.

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	4.21e-02	1.6e-02	2.557	0.0119 *
Block	-2.5e-04	6.1e-03	-0.041	0.9674
Agenttypeunfair	-1.3e-04	1.8e-02	-0.007	0.9944
Block:Agenttypeunfair	1.07e-02	8.6e-03	1.23	0.2191

Figure 10: The correlation between score and lie use



Next to the final offers and the favors we also look at the amount that players call the agent on their lying. We do this once more by plotting the amount of calling someone on a lie against the blocks and looking at a generalized linear mixed model along with some correlations between score and lie calling used. We do this only for the stupid and the unfair agent because we have gotten no data on the fair agent.

We see no significant effect from either block or agent type on the amount people call the agent on a lie. There seems to be some variance as can be seen in the plot, but this variance is not big enough.

There is also no correlation between the amount the subjects call the agent on a lie and the score. The values of these correlations are p-value = 0.8399 for the unfair agent and a p-value = 0.9657 for the stupid agent.

3.6. Saying something about MNS value

Figure 11: The use of saying something about MNS value

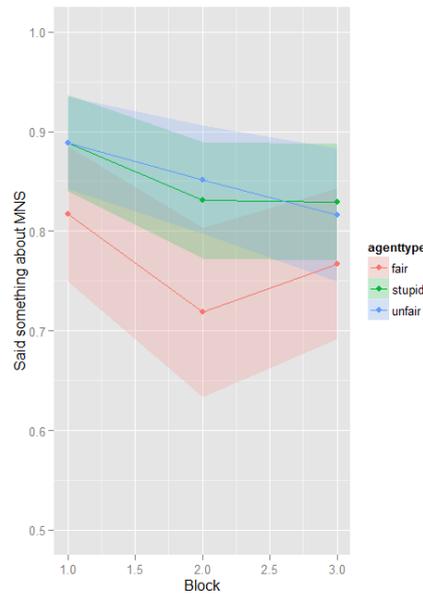


Table 7: Linear mixed model fit of block and agent type on saying something about MNS.

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	0.82602	0.0528	15.64	< 2e-16 ***
Block	-0.01316	0.01349	-0.97	0.32937
Agenttypestupid	0.12427	0.0412	3.016	0.00259 **
Agenttypeunfair	0.12865	0.0412	3.122	0.00182 **
Block:Agenttypestupid	-0.01974	0.01907	-1.03	0.30089
Block:Agenttypeunfair	-0.01535	0.01907	-0.80	0.42101

Furthermore we have looked at how many times a subject says something about his MNS value compared to when he does not say anything. We have done this with a plot of the amount of times a subject said something about his MNS against the block, along with a generalized linear mixed model of block and agent type on saying something about the MNS value.

We see here that there is a significant effect from both the stupid and the unfair agent on the amount people say something about their MNS value.

4. Discussion

In the results section, a lot of findings were given which will be discussed here. Only the major findings and the things that we were investigating are revealed in great detail. Some other, less substantial, findings are mentioned in section 4.6.

4.1. Score

First of all we see a great difference in the average score between on one hand the fair and unfair agent and on the other hand the stupid agent in figure 2. It seems that the stupid agent is not as hard to crack as we thought it would be, leading to a lot of participants getting a great sum of points out of this agent. After looking closely at the stupid agent, it seems that the amount that it is asking for is a little above a fair offer. This is good, but because it does also accept offers below what it is asking for and on top of that also gives favors and can get caught on lying, causes it to accept scores below the fair offer sometimes. The fact that the participants are able to fully exploit his weaker strategy does go to show that they do form some kind of mental image of that particular agent and that they adapt strategies based on opponent, something we wanted to find out. We know that they used a different strategy for the stupid agent, because if players played the same against the stupid agent as they would do against the fair or unfair agent, their score would not increase.

What we do not see is an influence of the block on the score, meaning that even though participants change strategies between agents, they do not learn to change their strategy in such

a fashion that they score more points later on in the experiment.

4.2. Lying about MNS

As can be seen in figure 3, there is a steady but certain decline in the amount participants are lying about their MNS value throughout the experiment. This also shows in table 3 with a significant influence of the block on the amount that people are lying about their MNS value. The correlations in figure 4 show that there is also a (almost) significant correlation between the amount that is lied to the unfair agent or the stupid agent and the score. This is not exactly what we expected, seeing as the stupid agent does not look at the MNS value of the opponent when determining whether or not to agree with a certain offer, but it does show that when people ask for a high offer, they back that up with a high MNS value.

Furthermore we see no effect from the agent type on the amount people lie about their MNS in the linear model, but it does look like people keep lying to the stupid agent more than to the others. This is to be expected since the stupid agent does not punish the player for having an aggressive and deceiving strategy.

4.3. Final offer use

We see in figure 5 that there is a big difference between the amounts that final offer is used between the different agents and that there is a shift in increase of usage for the fair agent and a decrease for the unfair agent. The generalized linear mixed model in table 4 however shows that only the influence of the unfair agent is statistically significant. If we look at the correlations shown in figure 6, we can see that there is a clear negative correlation between the score and the amount of final offers used against the fair agent. This is what it should be, meaning that the increase in final offer usage between block 2 and 3 for the fair agent is not something the participants should have incorporated into their strategies, but for some reason did. This might be a sign that in the end people were losing interest due to the experiment taking some time by then, and that the final offer use gave them a nice way to be done with each trial more quickly. We don't, however, see this across the other agents meaning that it might have

some other reason, for example having the feel that the fair agent who played nice was not punishing them too much for continuously making a final offer.

4.4. Favor use

When we look at figure 7 it is clear that people caught onto the effects of the favor use across the different agents. We see a high increase in the stupid agent across the blocks, which is also found in the general linear mixed model shown in table 5, with both the effect of the agent as well as the effect of the agent over the blocks being significant. This clearly goes to show that people found out that asking for favors with the stupid agent was a good way to get points, something they already found in the fair agent in the first block. There is a slight, but not significant, correlation with the score as can be seen in figure 8, telling us that a lot of participants got the first favor but never returned it. This is something that can be changed in further experiments, giving the agents themselves the ability to directly ask for a favor, instead of waiting for the player to make the first move.

4.5. Lie use

As told already in the methods section, something went wrong with the programming, meaning that there is no data for the usage of the lie button for the fair agent, since that agent did not use this in any of the calculations. The players do seem to call the unfair agent on the lying more (figure 9), but this is not a significant effect as can be seen in table 6. The amount that people do call someone on their lying is pretty low, which is something that was to be expected, since calling someone a liar is a big step and something not seen too often in the experiment of Stevens (unpublished) as well. The correlation however shows something weird. The correlation should have been a positive one, but there seems to be none present. When looking at the individual data we found that people were either using it relatively often, going over the limit where it would be of use, or pretty rarely. This is something that has to be reconsidered for further experiments.

4.6. Saying something about MNS

When looking at figure 11, we directly see a big difference in the amount that participants say something about their MNS values when playing against different agents, with both the unfair and the stupid agent having a significant effect on the amount participants say something about their MNS value when compared to the fair agent. This was not what we expected, since the effect it has on the agents is the same across all agent. When the subject says something about his MNS value, this changes 2 things in the agent. The first is that the agent will tell his own MNS value, giving the player a bit more information, and the second is that a different equation is used for determining a fair offer, seeing as the agent also has more information about the player. This last part of course is not so for the stupid agent. In the end there is no real way of saying whether saying something about your MNS is beneficial for a person, because this really depends on what they did in the previous turns. If a player said something about their MNS value and lied a lot about it, it might be beneficial for that particular turn to not say something about your MNS value, because then the other equation is used to determine whether or not the player is lying. In the end this might have an effect on score and strategy, but because players will in the end always pay for having an aggressive strategy, whether they try to change the equations used or not, this will only be a minor one.

4.7. Further findings

There was one major finding that was not expected. When looking at the data of the subjects per separate block, we found that the scores of people with different orders in which they faced the agents were rather different. People who started with the stupid agent had a very high score for that particular agent, but scored dramatically on the fair and unfair agent that followed after. Seeing as we only had 3 test subjects for the order stupid-fair-unfair and 3 test subjects for the order stupid-unfair-fair, this is something that could still be random chance, but it does seem that due to the different nature of the stupid agent, people don't get how the game normally works. They seem to quickly be able to see through the concept of the stupid agent, being able to pressure him into deals that are

only slightly beneficial for the agent and very beneficial for the player, but have no game plan once they face the fair and unfair agent and are not able to recover from that. This is something that should be looked into further, seeing as it possible that this means that people are only able to adapt to a certain extent once they have created an initial strategy.

Furthermore when talking with the subjects after the experiment, some were saying that the amount of output that the agent gave based upon the behavior of the player was not enough and that that made it hard for people to really feel like they were playing a proper opponent. It might be good to see whether it is possible to enhance this in further experiments as to create an even more realistic simulation with which the participant can interact.

5. Conclusion

Looking at the findings of the different options players had against the different agents it is safe to say that people do adapt a different strategy if they are playing different agents. Seeing as the favors that were added for this experiment have a definite impact on the play style of the participants we can conclude that there is an effect of adding more, realistic, options. The findings show that people have some sort of image of their opponent and are able to switch between strategies that look to work best for them.

However, we see that participants don't get better in playing the game, that is, score more points as they progress through the blocks. It is therefore uncertain what people see as a fitting strategy against a certain opponent, seeing as in a bargaining game it is the goal to get as much points as possible. It is possible that in these kind of situations people go for a score that is high enough, but also not too high, making them feel at ease with how much pressure they are putting on their opponent.

References

- Kelley, H. H., Beckman, L. L., & Fischer, C. S. (1967). Negotiating the division of a reward under incomplete information. *Journal of Experimental Social Psychology*, 3(4), 361-398.
- Liebert, R. M., Smith, W. P., Hill, J. H., & Keiffer, M. (1968). The effects of information and

magnitude of initial offer on interpersonal negotiation. *Journal of Experimental Social Psychology*, 4(4), 431-441.

- Schoeninger, D. W., Wood, W. D., (1969). Comparison of Married and Ad Hoc Mixed-Sex Dyads Negotiating the Division of a Reward. *Journal of Experimental Social Psychology*, 5, 483-499.

- Mascarenhas, S., Marques, N., Campos, J., & Paiva, A. (2013, December) A Model of Social Dynamics for Social Intelligent Agents. In *2013 AAAI Fall Symposium Series*.

- Daamen, J., (2013). Game of Nines: An exploratory study on the human-agent bargaining process. First year research project. Institute for Artificial Intelligence, University of Groningen

- Stevens, C., Game of Nines person vs person experiment. (unpublished). Institute for Artificial Intelligence, University of Groningen