metalogue

METALOGUE: Deliverable D2.4

Agent Models of the Game of Nines

The METALOGUE Consortium

March 2015

Version N°: [Version 1]

Main Authors: Christopher Stevens, Niels Taatgen, Fokie Cnossen

Project funded by the European Community under the
Seventh Framework Programme for
Research and Technological Development

Project ref. no.	ICT – 611073
Project title	METALOGUE – Multiperspective Multimodal Dialogue: dialogue system with metacognitive abilities

Document status	**[Draft, v1.0]**
Contractual date of delivery	December 2014
Actual date of delivery	February 2015
Document number	DX._
Deliverable title	Agent Models of the Game of Nines
Dissemination level	[Public]
Type	[Report]
Number of pages	[18]
WP contributing to the deliverable	WP 2
WP / Task responsible	University of Groningen
Contributing partners	University of Groningen
Reviewers	Internal Reviewer: Thomas Kleinbauer
Author(s)	Christopher Stevens, Niels Taatgen, Fokie Cnossen
EC Project Officer	Pierre-Paul Sondag
Keywords	Negotiation, Theory of Mind, Metacognition

Table of Contents

1. Executive Summary

Here we present a set of agents based in the ACT-R cognitive architecture that can play the game of nines with a human player. The agents implement strategies commonly discussed in the distributive negotiation literature. One agent uses a cooperative strategy (softline). The second agent uses an aggressive strategy (hardline). The third agent is a metacognitive agent that monitors the behavior of the opponent to determine the strategy he or she is using. It will then respond with the appropriate strategy.

These agents could be incorporated into a tutoring module that can be used to train students in distributive negotiation. We hypothesize that training against a metacognitive agent will help students learn to adapt to their opponents and also reflect on how their negotiation strategies are perceived by their opponents.

2. Introduction

In this deliverable we present a set of agents that function as virtual opponents for students learning distributive bargaining (see below). These agents embody strategies commonly employed in a variety of distributive negotiation tasks. The specific task used here is the Game of Nines, a two-person bargaining game developed by Kelley, Beckman, and Fischer (1967; See Deliverable 2.2 for a detailed description of the task.). The basic rules of the game are as follows. Two players must decide how to divide up nine points between them. Each player has a reserve price (Minimum Necessary Share or MNS) they must achieve in order to not lose points on the deal. Once the players agree on a point distribution, each player receives their allotted points minus their MNS. If the players fail to reach an agreement in a given round, they neither receive nor lose any points. Each player's knows their own MNS but not the opponent's. Players are free to share (or lie about) their MNS with the other player. All three agents are developed in an empirically-grounded cognitive architecture (ACT-R; Anderson et al., 2004). We expect that these agents will be useful as training partners, helping students learn how to adapt to different opponents and reflect on their own behaviors.

2.1 Distributive Bargaining Strategies

Distributive bargaining refers to a negotiation situation between two parties in which a fixed sum of goods or resources must be divided (Bartos, 1995). These bargaining situations are mixed motive because both parties are attempting to maximize their individual gains, but they gain nothing if they fail to agree. As a result, it is important to consider the other party's preferences when proposing a deal. In a situation like this, a negotiator could have one of two types of goals: cooperative or aggressive (Esser & Komorita, 1975; Van Kleef, De Dreu, & Manstead, 2006). A cooperative negotiator aims to maximize agreement and joint gains without sacrificing too much in terms of individual gains. An aggressive negotiator seeks only to maximize his/her individual gains. They consider the gains of their opponents only insofar as it allows agreements to be reached.

The cooperative and aggressive goal orientations described above are embodied in softline and hardline negotiation strategies respectively. These two strategy types have been studied intensively in the negotiation literature (Allen, Donohue, & Stewart, 1990; Huffmeier, Freund, Zerres, Backhaus, & Hertel, 2014; Osgood, 1962; Siegel & Fouraker, 1960). Distributive negotiation strategies are typically defined with respect to two primary features: opening bids and concession rates (Chertkoff & Conley, 1967). Softline strategies are defined by moderate opening bids and moderate concession rates. By contrast, hardline strategies are associated with extreme opening bids and low rates of concession (Huffmeier et al., 2014).

The aim of hardline negotiation strategies is to reduce the expectations of the opponent (Siegel & Fouraker, 1960). Making tough offers and refusing to concede communicate to the opponent that he/she should expect a less favourable deal from the exchange. This strategy forces the opponent into a dilemma. Accepting a hardline bargainer's terms may result in a gain for the opponent, but this will result in an unfair deal, giving a disproportionate amount of value to the hardline bargainer. Many studies have found that people often give in to this strategy, as this strategy usually results in better short-term economic gains (Allen et al., 1990; Huffmeier et al., 2014; Siegel & Fouraker, 1960). However, using aggressive strategies will usually cause the opponent to develop negative emotions toward the negotiator and the final agreement(s). This can harm cooperation in future negotiations, resulting in weaker economic gains (Huffmeier et al., 2014; Maaravi, Pazy, & Ganzach, 2014)

By contrast, softline negotiators attempt to establish a rapport with their opponents (Allen et al., 1990; Osgood, 1962). Softline negotiators give more generous offers and offer more concessions in the hope that their opponents will do the same. These strategies have been shown to result in better emotional outcomes and negotiation satisfaction in opponents (Huffmeier et al., 2014; Maaravi et al., 2014). Furthermore, they can even result in better economic outcomes than hardline strategies when they perfectly reciprocate the concessions of the opponent (Huffmeier et al., 2014). Unfortunately, these strategies can be exploited by opponents, resulting in worse economic outcomes (Allen et al., 1990; Huffmeier et al., 2014; Siegel & Fouraker, 1960).

2.2 The Role of Metacognition in Negotiation

Given that negotiation is a complex task with multiple strategies available, adapting to one's opponent can be challenging. A strategy that works well on one opponent may not work on another. For example, softline strategies are appropriate when others are generally willing to cooperate and interested in achieving positive joint outcomes. These strategies can lead to better joint outcomes for negotiators and will benefit their long-term relationship, resulting in better negotiations and agreements in the future (Curhan, Elfenbein, & Eisenkraft, 2010; Huffmeier et al., 2014). However, softline strategies can be exploited by hardline strategies. This basic tension is known as the "negotiator's dilemma" (Lax & Sebenius, 1986).

An important metacognitive skill in negotiation is the ability to read an opponent and determine an appropriate counter strategy. This type of monitoring requires knowledge of the types of strategies typically used in negotiation. As discussed above, the two main strategies identified in the literature are hardline and softline bargaining. An ideal negotiator should be able to quickly recognize both strategies and respond appropriately. Therefore, the metacognitive agent we present here has the ability to recognize behaviors commonly associated with hardline and softline strategies.

We hypothesize that playing against a metacognitive agent should be a useful exercise for students learning to negotiate. Because the agent adapts itself to the player's behavior, the player should consider the agent's representation of the player's strategy. A move that would make sense against a fixed-strategy agent may be ill-advised against an agent that will change strategies in response to the move. This may help to develop an ability known as second-order theory of mind. Second-order theory of mind refers to the ability to recognize that the opponent is reasoning about the player and to use this knowledge to shape the opponent's beliefs about the player (de Weerd, Verbrugge, & Verheij, 2013). The metacognitive agent presented in this deliverable not only attempts to infer the strategy of the user, but it also provides real-time feedback to the user about its evaluation of their strategy. We are currently running a study in which students train by playing against this agent with feedback. We hypothesize that this will result in an improvement in negotiation skill even when feedback is removed.

3. Description of Agents

3.1 General Mechanisms

The agents presented here are all based in the ACT-R 6.0 cognitive architecture (Anderson et al., 2004). The agents represent instance-based decision making, in which decisions are based on a collection of prior episodes (i.e. instances); it has been successful in modeling human behavior in a variety of decision-making tasks and 2-person games (Gonzalez & Lebiere, 2005; Lebiere, Wallach, & Taatgen, 1998; Logan, 1988). The strategies used by the agents, along with the knowledge needed for metacognitive monitoring, are represented as chunks in declarative memory. Declarative knowledge represents the past negotiation experiences of the agent as well as its current knowledge of the opponent. This knowledge is used to decide which strategy the agent should use as well as what the agent's next move should be. For example, a particular move might be represented by the following chunk:

Move1
My-Strategy: Cooperative
Value-of-my-current-offer: 3
Opponent's last move: Concede by 2
My-move: Concede by 2

This chunk tells the agent to lower its current offer by two in a particular situation. In this case, the situation is that the agent is using the cooperative strategy, its current offer is worth three points (after subtracting the MNS value), and the opponent just lowered their offer by 2.

Instance-based learning decision-making is analogous to reinforcement learning techniques in artificial intelligence (Kaelbling, Littman, & Moore, 1996). Like reinforcement learning, instances that result in better outcomes are more likely to be retrieved later. This is usually accomplished in one of two ways. Chunks may have their activation levels strengthened (Lebiere, Wallach, & West, 2000), or the results of applying the instance may be stored in the instance itself. In the latter case, these outcomes are used to estimate the expected value of the instance, and generally the instance with the best expected value will be chosen (Gonzalez & Lebiere, 2005).

One advantage of an instance-based approach is that it allows for a great deal of flexibility and transparency in the agents. To create a unique opponent, one needs only to change the instances stored in declarative memory. These instances need not represent all possible situations. If an instance cannot be found in memory that perfectly matches the current situation, ACT-R's partial matching mechanism is used to select the one that matches best. This mechanism is described in greater detail in Deliverable 2.2.

ACT-R estimates the importance of a particular instance using a quantity known as activation. The more recently an instance occurred, and the more frequently it is recalled, the more active that instance is. See Deliverable 2.2 for a more detailed description of how activation is computed. In the current agents, each instance represents a move that the agent should perform and the conditions under which the agent should perform it. These instances are presumably learned by the negotiator either through practice or prior instruction about how to negotiate. We do not model the process through which these instances are learned here.

We designed the agents to choose moves by retrieving instances from declarative memory. These instances correspond to the following moves: concede, insist, quit, and final-offer (See Filzmoser & Vetschera, 2008 for a discussion of these move types.). When concede is selected, the model reduces its current offer by one or two. When insist is selected, it resubmits its current offer. When final-offer is selected, it resubmits its current offer and indicates that it will not accept anything less. Each instance has slots representing the conditions under which it should be retrieved. This includes the strategy (cooperative, aggressive, neutral), the distance between the agent's current offer and its MNS value, and the opponent's previous move (both type of move and how far they reduced their offer). An agent is more likely to retrieve an instance if it matches its current strategy. Neutral instances are slightly less likely to be selected than those matching the current strategy, but are more likely to be selected than those matching the opposite strategy.

Three different agents are included here. The cooperative agent uses a softline bargaining strategy, emphasizing agreement over individual gain. The aggressive agent uses a hardline strategy, attempting to score as high as possible. Finally, the metacognitive agent flexibly switches between both types of strategy depending on the current opponent. It is important to note that the instances used by all three agents are not designed to cover every possible situation. Nor do they represent ideal cooperative or aggressive strategies. Rather, they provide examples that the agents can apply to the current situation through partial matching (see Deliverable 2.2).

3.1.1 The Agent Interface

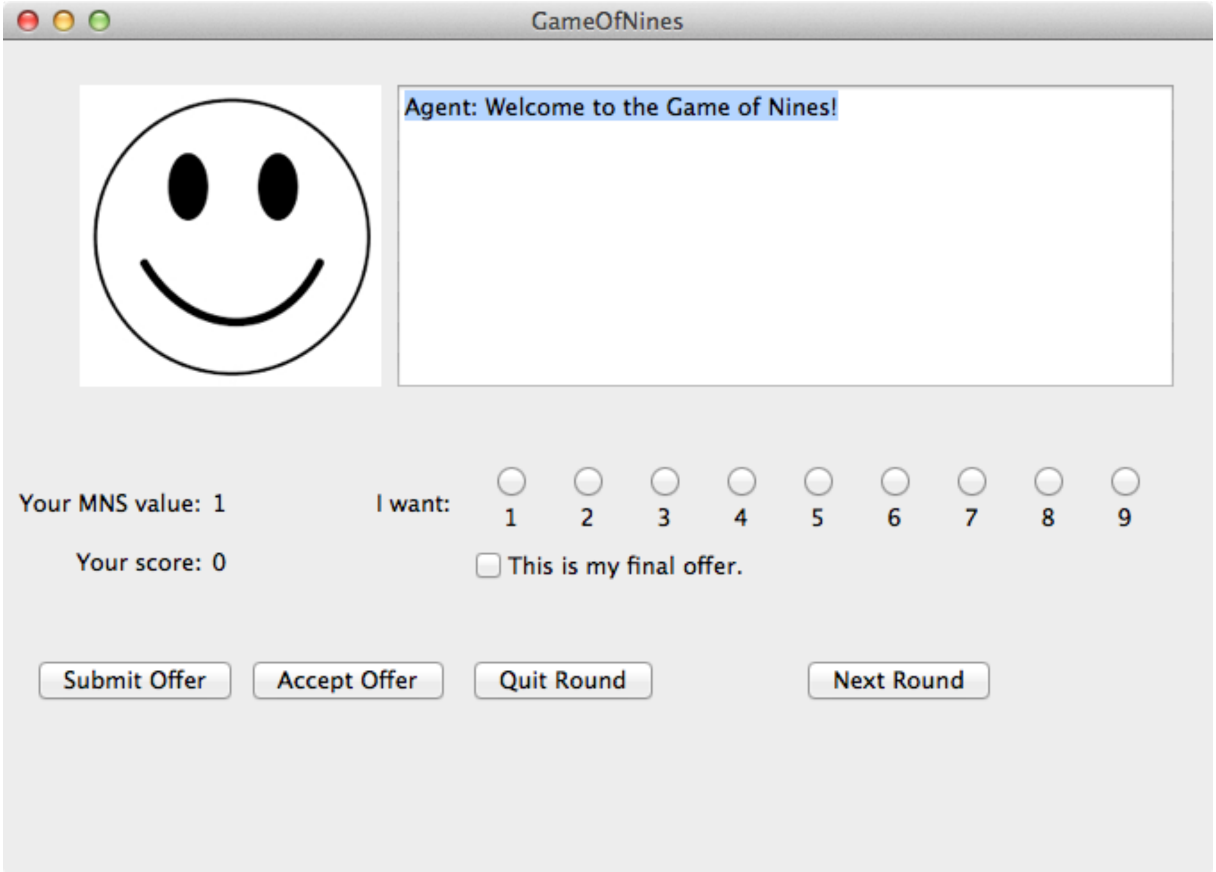

In order to run the agents, clozure common lisp (CCL) running on a Macintosh computer is required. The human player communicates with the agents through a window. The window contains a set of radio buttons labelled 1-9, corresponding to each of the possible offers a player can make. Once the participant selects their offer, they click the "Submit Offer" button to submit it to the agent. To make a final-offer, the player selects the corresponding checkbox before making their offer. The player also has the option of accepting the agent's offer and quitting without making a deal. After the player submits an offer, the agent examines the offer, selects its move, and communicates its move to the player through the chat box on the right. The metacognitive agent can also communicate to the player which type of strategy it is using (cooperative or aggressive) by changing the expression on the smiley face in the left window. A smiley face indicates that the agent believes the player is being cooperative (and that the agent will therefore respond cooperatively). An angry face indicates that the agent believes the player is being aggressive, and will therefore respond with an aggressive strategy. This will be useful for training purposes because it will provide real-time feedback about how the player's actions are being perceived by the agent. The player's MNS value and current score are displayed on the left side of the window.

The structure of the Game Of Nines task can be customized to fit the experimenter's needs. The lisp code running the experiment can be modified to allow any number of trials, and MNS values can be customized as well. All of these variables can be found at the head of the application file, and can be modified in any text editor. By default, the agent uses the same MNS values as in Kelley's (1967) Experiment 2. However, this can be modified to suit the training situation.

3.2 The Cooperative Strategy

The cooperative strategy is designed to maximize agreement and between the agent and the player. With respect to opening offers, the cooperative strategy is designed to give opening bids that are relatively moderate. A moderate opening offer yields a more even distribution of value, although it should still be more valuable to the proposer to allow room for concession (Huffmeier et al., 2014). The agent's instances provide several examples of opening offers like the following:

Initial-offer-1
Strategy: cooperative
Your-mns: 4
Your-opponent's-offer: 6
You-should-offer: 6

Thus the opening offer depends on the agent's mns value and the opponent's offer (if the opponent has made one). The extremeness of an opening bid is dependent upon the situational context and the bargaining range. The instances in the model provide examples of opening offers that give modest gains given certain mns values. Other mns values can be accommodated through partial matching (see Deliverable 2.2). Please see Appendix A for the instances used in the current implementation.

The cooperative concession strategy was designed to be one of reciprocity. When the opponent concedes, the cooperative instances indicate that the agent should concede by the same amount (up to 2 points). However, the decision to concede is also based on the distance between the agent's current offer and its mns value. The cooperative strategy eschews final offers, and will tend to quit if an opponent uses them. Appendix B displays the instances used for this decision.

3.3 The Aggressive Strategy

In contrast to the cooperative strategy, the aggressive strategy is designed to give more extreme opening bids and concede less often. Agents using this strategy will select higher opening offers on average (see Appendix A). Once again, these opening bids are based on the agent's mns value and the opponent's offer. The aggressive strategy is also more resistant to making concessions than the cooperative strategy and does not exhibit perfect reciprocity (see Appendix B). Finally, unlike the cooperative strategy, an agent using the aggressive strategy will accept a final offer if it can gain points from it. It also becomes more likely to use them as its own offer gets closer to its MNS value.

3.4 The Neutral Strategy

The neutral strategy is designed to be a middle-ground between the cooperative and aggressive strategies. Unlike the cooperative strategy, it will accept and use final offers. It is also less aggressive in its offers and final offers than the aggressive strategy. It serves two roles in the current agents. First, both cooperative and aggressive agents can use neutral moves. This helps prevent their behavior from becoming too rigid. For the metacognitive agent, the inclusion of a neutral strategy is important because it allows the agent to recognize when a move is ambiguous and does not clearly indicate one strategy over another. Instances guiding this strategy can also be found in the appendices.

3.5 The Metacognitive Agent

Like the cooperative agent, the metacognitive agent was designed to maximize cooperation with opponents. However, the metacognitive agent takes an approach of cautious cooperation. It evaluates each move made by the opponent as cooperative, aggressive, or neither. Then, it will choose the same strategy it believes its opponent is using. A logic similar to the tit-for-tat strategy in the prisoner's dilemma is applied here (Axelrod, 1980). When the agent detects a cooperative opponent, it will cooperate as well. When the agent detects an aggressive opponent, it will respond with an aggressive strategy to avoid exploitation. This meta-strategy has been successful in modeling human behavior in 2-person games (Gonzalez, Ben-asher, Martin, & Dutt, 2014; Juvina, Lebiere, & Gonzalez, 2014; Stevens, Taatgen, & Cnossen, 2015).

The metacognitive agent evaluates each move made by the opponent. The agent uses the same declarative knowledge to select its own moves and to evaluate the opponent's moves. When it observes an opponent's move it retrieves an instance from memory based on the move and the current context. The only difference is that instead of using its own context, the agent takes the perspective of the opponent. In other words, when the agent selects its own move, it starts with a strategy and looks up a move. When evaluating the opponent, it starts with the opponent's move and looks up a strategy. Once the instance is retrieved, the agent checks which strategy is associated with the instance (cooperative, aggressive, or neither). Both strategies are represented as declarative memory chunks. After an evaluation is made, the activation of the chunk corresponding to the appropriate strategy is increased. The model then determines its own strategy by selecting the most active strategy from memory.

The agent does not have access to its opponent's MNS values. However, it remembers its own MNS values and thus can infer the opponent's average MNS value (because on average both players will have the same MNS values). To do this, the agent stores chunks representing its own MNS value in memory, then it uses ACT-R's blending module to obtain a weighted average (Lebiere, 1999). The weight is determined by the activation level of the instance. When evaluating the opponent's move, the agent will use this average MNS value instead of its own current MNS value.

Before every move, the agent decides whether to use a cooperative or aggressive strategy by retrieving the most active strategy chunk. This strategy chunk is then used to select the appropriate instances from declarative memory. As an opponent shows a stronger tendency to cooperate or compete, the agent will become biased toward one of these two strategies. If a player does not show a tendency to use one strategy over the other, then the agent will select its strategy at random. The strategy-specific instances are the same as those used in the cooperative and aggressive agents.

4. Summary

Two important metacognitive skills for negotiators are the ability to adapt one's strategy to suit an opponent, and the ability to anticipate how one's actions will be interpreted by an opponent. We expect these agents will provide useful training partners for students to learn both of these vital skills. The single-strategy agents embody two prevailing distributive bargaining strategies, allowing students to gain experience dealing with these strategies. In addition, the metacognitive agent actively evaluates the player's actions and adjusts its behavior accordingly. This may help students learn about how their actions are perceived by opponents, encouraging them to consider not only their own perspective, but also their opponent's perspective in real negotiations.

5. References

Allen, M., Donohue, W., & Stewart, B. (1990). Comparing hardline and softline bargaining strategies in zero-sum situations using meta-analysis. *Theory and Research in Conflict Management*, 86–103.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*(4), 1036–60. doi:10.1037/0033-295X.111.4.1036

Axelrod, R. (1980). More Effective Choice in the Prisoner's Dilemma. *Journal of Conflict Resolution*, *24*(3), 379–403. doi:10.1177/002200278002400301

Bartos, O. J. (1995). *Modeling Distributive and Integrative Negotiations*. *Annals of the American Academy of Political and Social Science*.

Chertkoff, J. M., & Conley, M. (1967). Opening Offer and Frequency of Concession As Bargaining Strategies. *Journal of Personality and Social Psychology*, *7*(2), 181–185. doi:10.1037/h0024997

Curhan, J. R., Elfenbein, H. A., & Eisenkraft, N. (2010). The objective value of subjective value: A multi-round negotiation study. *Journal of Applied Social Psychology*, *40*, 690–709. doi:10.1111/j.1559-1816.2010.00593.x

De Weerd, H., Verbrugge, R., & Verheij, B. (2013). Higher-order theory of mind in negotiations under incomplete information. In *PRIMA 2013: Principles and Practice of Multi-Agent Systems*.

Esser, J. K., & Komorita, S. S. (1975). Reciprocity and concession making in bargaining. *Journal of Personality and Social Psychology*, *31*(5), 864–872. doi:10.1037/h0076692

Filzmoser, M., & Vetschera, R. (2008). A Classification of Bargaining Steps and their Impact on Negotiation Outcomes. *Group Decision and Negotiation*, *17*(5), 421–443. doi:10.1007/s10726-008-9106-1

Gonzalez, C., Ben-asher, N., Martin, J. M., & Dutt, V. (2014). A Cognitive Model of Dynamic Cooperation with Varied Interdependency Information Cleotilde, 1–42. doi:10.1111/cogs.12170

Gonzalez, C., & Lebiere, C. (2005). Instance-Based Cognitive Models of Decision- Making. In D. Zizzo & A. Courakis (Eds.), *Transfer of Knowledge in Economic Decision Making*. New York: Palgrave McMillan.

Huffmeier, J., Freund, P. a., Zerres, A., Backhaus, K., & Hertel, G. (2014). 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. doi:10.1177/0149206311423788

Juvina, I., Lebiere, C., & Gonzalez, C. (2014). Modeling trust dynamics in strategic interaction. *Journal of Applied Research in Memory and Cognition.*, 1–11. doi:10.1016/j.jmb.2005.03.040

Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement Learning : A Survey, 237–285.

Kelley, H., Beckman, L., & Fischer, C. (1967). Negotiating the division of a reward under incomplete information. *Journal of Experimental Social …*, *361398*. Retrieved from http://www.sciencedirect.com/science/article/pii/0022103167900042

Lax, D., & Sebenius, J. (1986). *The Manager as Negotiator*. Boston: Harvard Business School Press.

Lebiere, C. (1999). Blending: An ACT-R mechanism for aggregate retrievals. In *Sixth Annual ACT-R Workshop at George Mason University.*

Lebiere, C., Wallach, D., & Taatgen, N. (1998). Implicit and explicit learning in CT -R. In *Proceedings of the Second European Conference on Cognitive Modelling*.

Lebiere, C., Wallach, D., & West, R. L. (2000). A memory-based account of the prisoner's dilemma and other 2x2 games. *… of International Conference on Cognitive Modeling*. Retrieved from http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:A+memory-based+account+of+the+prisoner's+dilemma+and+other+2x2+games#0

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*. doi:10.1037/0033-295X.95.4.492

Maaravi, Y., Pazy, A., & Ganzach, Y. (2014). Winning a battle but losing the war : On the drawbacks of using the anchoring tactic in distributive negotiations. *Judgment and Decision Making*, *9*(6), 548–557.

Osgood, C. E. (1962). *An alternative to war or surrender*. Urbana: University of Illinois Press.

Siegel, S., & Fouraker, L. E. (1960). *Bargaining and group decision making: Experiments in bilateral monopoly. Journal of Urban Economics* (Vol. 19). Retrieved from http://psycnet.apa.org/psycinfo/1961-01746-000

Stevens, C. A., Taatgen, N. A., & Cnossen, F. (2015). Metacognition in the Prisoner's Dilemma. In N. A. Taatgen, M. Van Vugt, J. Borst, & K. Melhorn (Eds.), *Proceedings of the 13th annual International Conference on Cognitive Modeling*. Groningen.

Van Kleef, G. a, De Dreu, C. K. W., & Manstead, A. S. R. (2006). Supplication and appeasement in conflict and negotiation: The interpersonal effects of disappointment, worry, guilt, and regret. *Journal of Personality and Social Psychology*, *91*(1), 124–42. doi:10.1037/0022-3514.91.1.124

6. Appendix A: Opening offer instances

Strategy	My MNS Value	Opponent's Opening Offer	My Opening Offer
Cooperative	1	6	5
	4	6	6
	1	5	5
	4	8	7
Neutral	3	7	7
Aggressive	4	8	9
	1	6	7
	4	7	8
	1	7	8

7. Appendix B: Concede/Insist Instances

Strategy	Net Value of My Offer	Opponent's Concession Amount	Then reduce offer by…
Cooperative	3	2	2
	2	2	2
	2	1	1
	1	1	1
	1	0	0
	3	0	1
Neutral	2	2	1
	3	1	1
	1	1	0
	2	0	0
	3	0	0
Aggressive	3	1	0
	2	1	0