metalogue

METALOGUE: Deliverable D2.3

Cognitive Model of a Multi-Issue Negotiation Task

The METALOGUE Consortium

[September] [2015]

Version N°: [Final]

Main Authors: Harmen de Weerd, Christopher Stevens, Fokie Cnossen, Niels Taatgen

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	**[Final version]**
Contractual date of delivery	_April _____ 2015
Actual date of delivery	_October 2015
Document number	D2.3
Deliverable title	Cognitive Model of a Multi-Issue Bargaining Task
Dissemination level	[Public]
Type	[Report]
Number of pages	[24]
WP contributing to the deliverable	WP 2
WP / Task responsible	University of Groningen
Contributing partners	
Reviewers	Olga Petukhova
Author(s)	Harmen de Weerd Christopher Stevens Fokie Cnossen Niels Taatgen
EC Project Officer	Pierre-Paul Sondag
Keywords	Metacognition, Cognitive Model, Negotiation

© METALOGUE Consortium [2015]

Table of Contents

© METALOGUE Consortium [2015]

1. Executive Summary

In this deliverable, we develop a new empirical paradigm in which to study metacognitive reasoning in negotiation. This task expands upon the Game of Nines, offering greater opportunity to examine integrative behaviour in negotiation. We also present an instance-based model capable of performing this task with other agents.

2. Introduction

Through the Game of Nines, we have investigated a particular bargaining setting in which two negotiators have to decide how to divide nine identical and indivisible goods among themselves. Such a task involves mostly *distributive bargaining* (Bartos, 1995; Fisher & Ury, 1981; Raiffa, Richardson, & Metcalfe, 2002; Raiffa, 1982) in the sense that negotiators engage mostly in competitive behavior.

In this document, we extend the Game of Nines to a multi-issue bargaining situation. In this setting, negotiators have preferences that are not necessarily competitive. This extends the Game of Nines setting to an *integrative bargaining* situation, where negotiators do not only decide on how to 'divide a pie' among themselves, but also have the opportunity to 'enlarge the pie' (Fisher & Ury, 1981; Raiffa et al., 2002; Raiffa, 1982) by leveraging differences in the relative importance of issues among negotiators. Our multi-issue bargaining setting is partially observable in the sense that negotiators know their own preferences, but not the preferences of their trading partner. This creates opportunities for a negotiator to cooperate with his trading partner in addition to competitive strategies.

To determine how human negotiators behave in this multi-issue bargaining setting, we report the findings of a pilot experiment. In addition, we develop an agent model for negotiation intended to describe and analyze the way human negotiators play the game.

In Section 3, we describe the multi-issue bargaining setting in more detail. Section 4 reports the results of the pilot experiment with human participants, and outlines general strategies observed in human negotiation. Section 6 introduces an agent model for a negotiating agent in the multi-issue bargaining setting.

3. Multi-issue bargaining setting

In the Game of Nines, two negotiators decide on how to divide nine indivisible and indistinguishable goods among themselves (Kelley, Beckman, & Fischer, 1967). That is, the Game of Nines involves two negotiators attempting to agree on one of 10 possible outcomes, over which they hold strictly competitive preferences. Whenever one of the negotiators prefers some outcome a_1 over a_2, the other negotiator prefers outcome a_2 over a_1. In addition, the Game of Nines has an opt-out outcome. A negotiator can decide unilaterally to withdraw from negotiation, in which case both negotiators end up with zero score. The addition of this opt-out adds a cooperative element to the Game of Nines in the sense that there exists at least one outcome a_3 that both negotiators prefer over no negotiation. However, since negotiating is costly in the Game of Nines, negotiators prefer the opt-out outcome over some of the possible negotiated outcomes.

In this document, we extend the Game of Nines to a multi-issue bargaining situation that allows arbitrary preferences over outcomes. That is, when one negotiator in the multi-issue bargaining situation prefers outcome a_1 over outcome a_2, this does not imply that the other negotiator prefers outcome a_2 over a_1. The other negotiator may be indifferent between the outcomes, or even have the same preference as his trading partner. This increases the opportunities for cooperative strategies that rely on maximizing the total value of the negotiated agreement (enlarging the pie) in addition to maximizing one's own value at the expense of the trading partner (dividing the pie).

The specific setting we consider involves a multi-issue bargaining scenario in which a representative of a city council and a representative of small business owners negotiate over the implementation of new anti-smoking regulations. The negotiation involves four issues, each with four or five different options (see Table 1). The task of the negotiators is to negotiate an agreement, which assigns exactly one option to each issue. Despite this simple setup, this setting allows for 400 different possible negotiation outcomes in addition to the opt-out outcome.

In our setup, each negotiator has preferences that assign a value to each possible option. Higher values are associated with more preferable options. The value of a negotiated outcome a is calculated as the sum of the values of the agreed-upon options in a. A negotiator therefore aims for a negotiated outcome that his preferences assign as high a value as possible.

Similar to the Game of Nines, negotiators may unilaterally decide to end negotiation. This opt-out outcome yields both negotiators zero score. A negotiator therefore considers any outcome a that his preferences assign a negative value to be unacceptable. In such a case, the negotiator would prefer to opt out of negotiation over accepting outcome a.

Table 1 : The multi-issue bargaining setting involves four separate issues with four or five options

1. **Scope of smoking ban**
    a. All outdoor smoking allowed
    b. No smoking in public transportation
    c. No smoking in public transportation and parks

d. No smoking in public transportation, parks, and open air events

2. Taxation of tobacco products

a. No change in tobacco taxes

b. 5% increase in tobacco taxes

c. 10% increase in tobacco taxes

d. 15% increase in tobacco taxes

e. 25% increase in tobacco taxes

3. Anti-smoking campaign

a. Flyer and billboard campaign in shopping district

b. Anti-smoking posters at all tobacco sales points

c. Anti-smoking television advertisements

d. Anti-smoking advertisements across all traditional mass media

4. Enforcement of tobacco sales restrictions to minors

a. Police fines for minors in possession of tobacco products

b. Ban on tobacco vending machines

c. Police fines for selling tobacco products to minors

d. Identification required for all tobacco purchases

e. Government-issued tobacco card for all tobacco purchases

© METALOGUE Consortium [2015]

4. Experimental setup

To determine the way human negotiators behave, we let participants play the multi-issue bargaining setting described in Section 1. Each experiment involves a pair of participants that perform a number of separate negotiation scenarios. At the start of the experiment, one of the participant is randomly assigned the role of city council representative. The other participant is assigned the role of the small business representative. Each participant then receives their cover story and instructions (see Annex 8.1 and Annex 8.2), as well as their preference profiles for each scenario. For each preference profile, each option was assigned one of nine possible values, which was communicated to the participant through colors, as shown in Figure 1. Brighter orange colors indicated increasingly more negative options, while brighter blue colors indicated increasingly more positive options. The use of color rather than numbers introduces a form of uncertainty in the exact value of a given agreement, which we argue to be closer to real-life negotiations.

Our experiment introduced minimal rules for the negotiation process. Participants were asked to negotiate for an agreement with as high as possible a value according to their preference information. They were not allowed to accept agreements that had a negative value, and participants were not allowed to show their preference information to their trading partner. We did not impose any further rules on the negotiation process. In particular, no time limit was set for negotiations, and participants were not required to follow guidelines for making offers. During the experiment, the dialogue between the participants was captured with two headset microphones to record the voices of the participants separately.

Figure 1: Values of issues can take on one of nine different values, presented to participants

5. Experimental results

In our pilot experiment, we let four pairs of participants play nine different scenarios (see Annex 8.3) of the multi-issue bargaining game. Each pair of participants negotiated for 75 minutes or until all nine scenarios had been resolved. A scenario was considered resolved when both participants had agreed on one option for each issue, or when one participant decided to end negotiations. In the following subsections, we describe the negotiation process, negotiated outcomes, and opportunities for improvement separately.

5.1 Process of negotiation

In automated negotiations, agents typically negotiate through a sequence of offers. These offers are commitments to a certain negotiation outcome. That is, the agent is bound to the offered negotiation outcome if the trading partner accepts. In participant negotiation, we see a different pattern of behavior. Participants rarely made a concrete offer as a binding commitment. Rather, participant actions were focused mainly on obtaining and providing preference information. Participants often provided their trading partner with (possibly inaccurate) information about their own preferences concerning one or more issues, or requested information about the preferences of their trading partner.

The most popular negotiation strategy among participants was to start by negotiating single issues separately. Participants typically alternated in suggesting one of the options of a given issue. If the negotiators succeeded in finding an option that was considered to be acceptable to both negotiating parties, this issue would become part of a tentative agreement. Participants typically avoided negotiating on multiple issues at once until no mutually acceptable option was found for a given single issue.

Importantly, the participants' suggestions are not commitments. Even when negotiators have tentatively agreed on each issue separately, a negotiator may still reject the resulting agreement. In some cases, negotiators used this to strategically choose their suggestions. For example, suppose a negotiator suggested a 10% tax on tobacco products, which would actually maximize the score of his trading partner for the taxation issue. Instead of accepting, however, the trading partner would sometimes suggest a different option (say, 15% tax) that would yield him a lower score. Later, this would allow the trading partner to suggest `conceding' to a 10% tax in return for a better option on, say, the anti-smoking campaign. That is, the trading partner would frame his suggestion as a concession without actually giving up anything.

Note that the strategic use of the asymmetry in information did not always yield the desired outcome. In one case in particular, confusion over the preferences of the trading partner lead to a suboptimal outcome. In this case, both participants believed they were making a concession to the position of their trading partner, which ended up in a worse outcome for both parties.

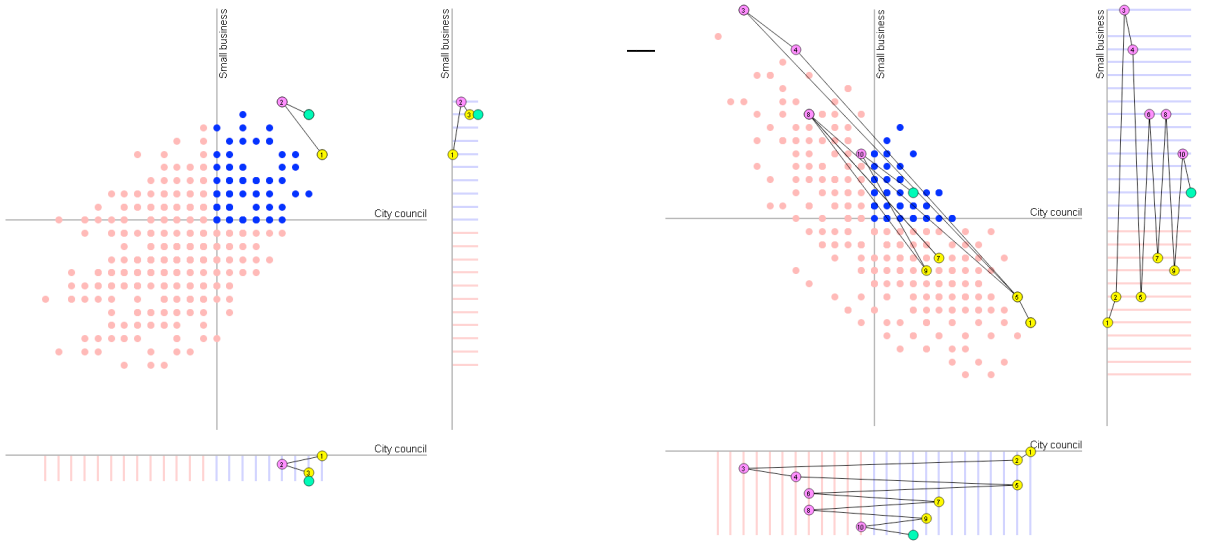

(a) Cooperative scenario (b) Competitive scenario

Figure 2: Negotiation space and negotiation outcome for (a) a cooperative scenario and (b) a competitive scenario. The top left panel in each figure shows the possible negotiation outcomes in terms of the score for the small business owners representative (vertical) and the city council representative (horizontal). Blue dots indicate outcomes that are acceptable to both negotiators, while red dots indicate unacceptable outcomes. Large, unnumbered dots show the best possible outcomes for the small business owners representative (purple) and the city council representative (yellow), as well as the final agreement (teal). The sequence of offers made during the negotiation is indicated by connected, numbered dots. The color of these dots shows which negotiator made the offer. The panels on the right and the bottom show the same information from the perspective of the small business owners representative and the city council representative, respectively. Other negotiation strategies exist, and one pair of participants followed such a different strategy. Instead of negotiating single issues separately, one of the negotiators would propose a draft agreement on all issues at once, after which the other negotiator would propose changes to the draft until a mutual agreement was reached.

5.2 Negotiated outcomes

In this subsection, we describe general patterns in the negotiated outcomes. Figure 2 shows examples of two negotiation processes, consisting of three panels. The top left panel shows the negotiation space for each scenario. Each dot shows a possible negotiation outcome in terms of the score of both negotiators. Dots that appear more towards the top indicate negotiation outcomes that are preferable to the representative of small business owners, while dots that appear more towards the right are more preferable for the representative of the city councils.

The distribution of dots across this negotiation space determines the type of negotiation. Figure 2a shows a positive correlation between the values of negotiated outcomes for both negotiators, indicating that the underlying scenario is cooperative. Similarly, Figure 2b shows a negative correlation, indicating a more competitive scenario. When no clear correlation exists between the valuations of the two negotiators, the scenario can be characterized as a neutral negotiation setting.

© METALOGUE Consortium [2015]

Blue dots indicate the zone of potential agreement, which are the negotiation outcomes that are acceptable to both negotiators. Red dots are outcomes that are unacceptable for at least one of the negotiators. The best possible outcomes for the representative of small business owners and for the representative of city councils negotiators are indicated by a larger, unnumbered dots (purple and yellow, respectively).

The actions of the negotiators are shown as a connected series of numbered negotiation outcomes corresponding to the sequence of offers that negotiators made. At the start of a negotiation, a suggestion or offer typically only covers a subset of the issues. In these cases, the options with the highest value for the negotiator making the offer have been selected for the issues that had not been discussed at that point in the negotiation. The final outcome of a negotiation is indicated as a teal dot.

In addition to the negotiation space, each graph in Figure 2 also shows the perspective of the representative of small business owners (top right panel) and the perspective of the representative of the city council (bottom left panel). These panels show that the perspectives of the two negotiators can be very different. For example, Figure 2b shows that from the perspective of the representative of the city council, offer 4 was a concession to the position of his trading partner, while the representative of small business owners would disagree.

A negotiated outcome is Pareto efficient if none of the negotiators could have achieved a higher score for themselves without a reduction in score of the other negotiator. The results show that in cooperative settings, negotiations tend to resolve more quickly and are more likely to be Pareto efficient. In competitive settings, on the other hand, negotiations tend to take more rounds of offers, and are more likely to result in an outcome that is not Pareto efficient.

5.3 Possibilities for negotiation improvement

Our results show several possible ways in which the negotiation strategies of participants could be improved. Firstly, despite the fact that participants were told that they were not allowed to accept a negative outcome, six out of the 28 negotiations (21%) resulted in an unacceptable outcome. Although this can be partially attributed to the color coding of option, which makes it more difficult for participants to judge the acceptability of a negotiated outcome, accepting an unacceptable offer also indicates that the negotiator could have negotiated for an outcome that would have yielded a significantly higher personal score.

In addition, 11 of the 28 negotiated outcomes were not Pareto efficient. That is, at least one of the negotiators could have increased his score without affecting the score of his trading partner. This shows that negotiators did not optimally explore the negotiation space. Indeed, some participants were more interested in reaching an acceptable outcome than an agreement that would yield them as high a personal score as possible. Negotiating with agents may help participants to explore the negotiation space more effectively.

An important metacognitive skill in negotiation is reasoning about opponent goals based on the information they give. Another important, though more difficult, skill is influencing your opponent's reasoning about your goals. At the most basic level, this means selectively communicating information about your own goals to your opponent. Although participants typically reasoned about the goals of their trading partner, they failed to let their trading partner know their

own preferred outcome. This may lead to situations in which both negotiators feel that they have conceded more to the position of their trading partner because neither of them is aware of the concessions that the other has made. Agents may help participants to incorporate metacognition in their negotiation behavior.

To summarize, agents could help negotiators to ensure

1. Acceptability of the negotiated outcome;
2. Sufficient exploration of the negotiation space; and
3. Effective use of metacognitive abilities.

6. The Cognitive Model

Our cognitive model, like our previous models of the Game of Nines (see deliverables 2.2 and 2.4) is based on the declarative memory mechanisms of the ACT-R architecture (Anderson, 2007) and Instance-Based Learning (IBL) theory (Gonzalez & Lebiere, 2005; Logan, 1988). Both of these theories have been reviewed extensively in previous deliverables (see Deliverables 2.2 and 2.4), so we won't elaborate on them further here. In the following section, we will describe the design of the model.

6.1 Simulated Environment

A virtual task environment was developed for the models in Java. Both the task and the model were developed in the Java programming language. The environment gives both models information about their own preferences and the actions of the opposing model. The models take turns making moves until one model quits or both have selected the same option on all dimensions. At the beginning of a turn, the model receives the opponent's current offer on all issues (if any) and any additional information (or requests for information) provided by the opponent. The model can then update its own offer on as many of the dimensions as it chooses and can choose to share (or request) information of its own. Preferences are represented using the numerical values from the Table 4 in Appendix 10.

6.2 Model Design

On every round, the model begins by reading in the opponent's positions on all dimensions and all extra information provided by the opponent. It compares this information to the opponent's position during the previous round to classify the opponent's move (see the section on actions for possible classifications). The model then classifies the opponent's move as cooperative, aggressive, or neutral by retrieving an instance from declarative memory. Next, the model decides whether the opponent is behaving cooperatively or aggressively by determining which strategy is more active in declarative memory. After identifying the opponent's strategy, the model will change its strategy to match the opponent and will select a move of its own.

In our behavioral data, it was common for participants to focus on one or two issues at a time. For this reason, we have designed the model to do the same. On every round, the model makes a decision regarding one particular dimension of the negotiation. The model selects the focal dimension by the following criteria:

1. If the opponent has made an information request regarding a certain dimension, the model will focus on that dimension.

2. If the opponent has made a concrete offer on a dimension, it will select that dimension. Otherwise, the model will randomly select an unresolved dimension.

3. If the negotiation appears to be deadlocked on a dimension and the opponent has been sufficiently cooperative, the model will attempt to make a concession on another unresolved dimension (see the trade off action below).

The model classifies the opponent as cooperative or aggressive using a similar method as that described in Deliverable 2.2. Every time the opponent makes a cooperative move, the activation level of the cooperative strategy is strengthened. Similarly, aggressive moves strengthen the aggressive strategy. To decide whether an opponent is cooperative or aggressive at a given time, the

model simply determines which strategy is more active, according to ACT-R's activation equation (see Deliverable 2.2). For the present model, we use a small amount of activation noise (s = .01).

The model also stores information about the preferences of both parties in its memory. The model starts with full information about its own preferences. When the opponent shares its preference on an option, the model will store that as well. This helps the model determine which options may be the most mutually satisfactory for both parties. The model also keeps track of the information it has shared with its opponent. This is important for metacognitive reasoning, because some actions may appear aggressive until one considers what the opponent knows about the model's preferences. For example, the opponent may select an option that has a high negative value for the model. If the opponent knows about the models preference, then the move is aggressive. But if it doesn't, the move might still be cooperative. For the purposes of the current simulation, we assume all of these memory representations to be permanent and error-free. Though it is certainly possible for a negotiator to misremember (or misunderstand) an opponent's preference, the phenomena of interest here do not require simulating these types of errors.

6.3 Negotiation Strategies

The metacognitive model has access to both cooperative and aggressive strategies. It employs IBL to identify the strategy of the opponent and adjust its own strategy accordingly. Specifically, it uses a cooperative strategy against cooperative opponents and an aggressive strategy against aggressive opponents. The rationale for this approach has been outlined in previous deliverables (2.1, 2.2, and 2.4). The relative characteristics of both of these strategies are outlined below.

The Cooperative Strategy. Cooperative negotiators will share information about their preferences with their opponents and will engage in problem-solving behaviors to attempt to find mutually beneficial agreements (De Dreu, Weingart, & Kwon, 2000). When the current model uses the cooperative strategy, it prefers the options that have the highest collective value for both parties. But if it does not have enough information to make this determination, it will seek out information from the opponent. In cases of deadlock, the cooperative strategy will not engage in positional bargaining tactics like anchoring or threatening. Instead, it will attempt to find dimensions where a trade-off is possible.

The Aggressive Strategy. Aggressive negotiators, like some of the participants in our behavioral data, will often withhold information from their opponents (De Dreu et al., 2000). These negotiators fear weakening their power in the negotiation by sharing too much. Therefore, these negotiators focus on asserting their own preferred positions rather than exploring the space of possible agreements (Fisher & Ury, 1981). In our model, the aggressive strategy is characterized by a focus on positional bargaining, and a general reluctance to share information. An agent using this strategy will rarely ask its opponent for preferences, and will usually ignore requests for information from the opponent. Instead, it finds its own ideal offer, state it, and insist upon it in hopes of making the opponent concede. In cases of deadlock, an aggressive agent will threaten to end the negotiation or make very small concessions. If the opponent makes a threat, the aggressive strategy will accept an offer only if it can gain a significant number of points from it.

6.4 Instance Design

Similar to our cognitive model of the Game of Nines (deliverables 2.2 and 2.4), our metacognitive model reasons about the intentions of its opponent by taking the opponent's perspective and comparing their behavior to cooperative and aggressive behaviors stored in the model's memory. These behaviors are stored as instances. An instance is a unit of memory that represents an action and a context in which to use that action. Our theory of memory is based on ACT-R's declarative memory module (Anderson, 2007). The declarative memory mechanisms contained in the model are built partially using code from a Java implementation of the ACT-R architecture (Salvucci, 2013).

The model has three types of instances: cooperative, aggressive, and neutral. The cooperative and aggressive instances describe behavior that is typical of their respective strategies according to the literature (Bottom & Studt, 1993; De Dreu et al., 2000; O'Connor & Arnold, 2001; Tinsley, O'Connor, & Sullivan, 2002). The instances are used both to implement the strategies and to identify them in opponents. Neutral instances describe behaviors that are not indicative of either strategy. The instances are hand-coded by the modelers in order to achieve a model that plays effectively and believably.

Instances are schematic representations, meaning that they consist of a set of slot-value pairs. The following are the possible slots in the instances in our model: They encode important information about the state of the negotiation (conditions) and responses to those states (actions).

6.4.1 Conditions

Strategy. This slot encodes the strategy associated with the instance (cooperative, aggressive, or neutral).

My-preferences (and opponent-preferences). There is a slot that encodes the model's preference score for each possible option (five in total) and a corresponding slot that codes the opponent's preference. This slot, combined with the *opponent-preference* slot, helps the model determine the relative importance to assign to each party's preferences. Higher numbers in the *my-preferences* slots tell the model to assign more importance to its own preferences than the opponent's.

Aggressive instances have high values in the my-preferences slots but don't specify the opponent's values, leading the model to ignore its opponent's interests. Cooperative instances have equal values in the my-preferences and opponent-preferences slots, specifying a more balanced approach.

A nil value in the opponent-preferences slot means the model does not yet know the opponent's preference for a given option. Some cooperative instances have nil values in these slots, telling the model when to request additional information (see the request-preference action in the next section).

Unresolved-dimensions. This slot indicates the number of dimensions in the negotiation for which no agreement has been reached. This helps the model determine whether it is possible to propose a trade off.

Turn. This slot records the number of offers that have been made in the current negotiation. This helps the model to know when a negotiation has reached a deadlock and it is necessary to

change tactics.

Opponent-action. This represents the move that the opponent has just made.

6.4.2 Actions

Select option. Make a concrete offer to select a specific option on a dimension. Aggressive instances will tell the model to select their most preferred options regardless of the opponent's preferences. Cooperative instances will attempt to select options that maximize both players' scores.

Request preference. Ask the opponent for their preference score on a given option. This action is associated with cooperative instances and is used when the model doesn't have enough information about the opponent's preferences to choose an option.

Share preference. Tell the opponent your preference score on a specific option. In the cooperative strategy, the model will share preferences when the opponent requests them. In the aggressive strategy, the model will tend to ignore the opponent's requests for information.

Threaten. Claim that you are going to end the negotiation if the opponent does not agree to a specific option. This option is associated with the aggressive strategy but not with the cooperative strategy.

Quit. End the negotiation without agreement.

Trade off. If the negotiation is deadlocked on the current dimension, search for another dimension on which you can change options. This move is more complex than the others, and requires the model to perform another search for a suitable option for another dimension.

6.4.3 An example instance

In this section we present an example of one instance found in the model and explain how its slots and data are used:

Name: coop6

Strategy: cooperative

My-Preference1: 4

Opponent-Preference1: nil

Turn: 0

My-move: Request-preference 1

This instance is designed to tell the model that it needs more information about the opponent's preferences for option 1. The strategy slot tells the model that this should be used when the model is in a cooperative state. The My-Preference1 slot contains a 4, which means this instance is more likely to apply when the model has a high preference score for option 1. The Opponent-Preference slot contains nil, which means the opponent has not yet shared its preference score with the model. Turn contains zero, meaning that this instance will be used more often in the early turns

of the negotiation. As the negotiation proceeds, the turn number will be farther from zero and this instance will be less likely to be retrieved. Finally, the my-move slot tells the model to request the opponent's preference value for option 1. So in plain language one can interpret this instance as "When you are using a cooperative strategy, and you strongly prefer option 1, but you don't know if your opponent likes option 1, ask your opponent what their preference value is for option 1."

6.5 Further Development

In this deliverable we describe the core mechanisms of our metacognitive model. This model serves two primary purposes in the greater context of the Metalogue system. First and foremost, this model will interact with the dialogue manager to evaluate ongoing negotiation and provide strategic feedback to learners. Later deliverables will describe in greater detail how this will be accomplished. Secondly, we envision a separate, strategic tutoring module with which learners can interact and receive feedback on their negotiation strategies independently from the feedback about verbal and body-language tactics that the rest of the system will provide. Strategic adaptation is a very important component skill in metacognition and this module will provide learners the opportunity to engage in focused practice on this skill. This will be especially helpful if the overall system identifies strategic adaptation as one of the student's areas of weakness.

6.6 Summary

The model we have presented here is designed to be a cognitively plausible model of human metacognition in the context of the multi-issue negotiation task. The key assumption is that strategic adaption and move selection are based on instances in declarative memory. The result is a model capable of adapting to its opponents, whether they use cooperative or aggressive tactics.

7. Discussion and conclusion

We have introduced a multi-issue bargaining extension of the Game of Nines, in which two players negotiate an agreement on one of 400 possible negotiation outcomes. Experimental results show interesting aspects of the actions that participants take while negotiating, in particular regarding the use of asymmetric information and strategic deception. These results suggest that participants are less likely to obtain a Pareto optimal outcome when the preferences of negotiators are negatively correlated. Especially in such cases, agents may be helpful in assisting participants in learning negotiation skills. In this deliverable, we have developed a candidate agent for this task. This agent is capable of metacognitive reasoning about its opponent's goals and intentions. In addition, the model exhibits behaviors that the human negotiators often lack (e.g. actively exchanging information about preferences). We expect that training with this model will encourage the use of metacognitive skills in negotiation.

References

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford University Press.

Bartos, O. J. (1995). Modeling distributive and integrative negotiations. *Annals of the American Academy of Political and Social Science*, *542*, 48–60.

Bottom, W. P., & Studt, A. (1993). Framing Effects and the Distributive Aspect of Integrative Bargaining. *Organizational Behavior and Human Decision Processes*, *56*(3), 459–474. doi:10.1006/obhd.1993.1064

De Dreu, C. K., Weingart, L. R., & Kwon, S. (2000). Influence of social motives on integrative negotiation: a meta-analytic review and test of two theories. *Journal of Personality and Social Psychology*, *78*(5), 889–905. doi:10.1037/0022-3514.78.5.889

Fisher, R., & Ury, W. L. (1981). *Getting to Yes: Negotiating Agreement Without Giving In*. London: Penguin Group.

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.

Kelley, H. H., Beckman, L., & Fischer, C. (1967). Negotiating the division of a reward under incomplete information. *Journal of Experimental Social Psychology*, *3*, 361–98.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492–527. doi:10.1037/0033-295X.95.4.492

O'Connor, K. M., & Arnold, J. A. (2001). Distributive spirals: negotiation impasses and the moderating role of disputant self-efficacy. *Organizational Behavior and Human Decision Processes*, *84*, 148–76. doi:10.1006/obhd.2000.2923

Raiffa, H. (1982). *The art and science of negotiation*. Cambridge: Belknap Press.

Raiffa, H., Richardson, J., & Metcalfe, D. (2002). *Negotiation Analysis: The Science and Art of Collaborative Decision Making*. Cambridge: Belknap Press.

Salvucci, D. D. (2013). Integration and reuse in cognitive skill acquisition. *Cognitive Science*, *37*(5), 829–860. doi:10.1111/cogs.12032

Tinsley, C. H., O'Connor, K. M., & Sullivan, B. a. (2002). Tough guys finish last: The perils of a distributive reputation. *Organizational Behavior and Human Decision Processes*, *88*(2), 621–642.

8. Annexes

8.1 City council representative cover story

In this experiment, you will play the role of a professional mediator employed by the government. In this capacity, your task is to negotiate the implementation of an anti-smoking campaign. Although the campaign is nation-wide, cities are allowed to make their own decisions concerning the details of the implementation. In fact, city councils turn out to be quite varied in their preferences. For this reason, you will enter separate negotiations with your trading partner for each city. During these negotiations, the following four aspects need to be decided on:

- Scope of smoking ban;
- Taxation of tobacco products;
- Anti-smoking campaign; and
- Enforcement of tobacco sales restrictions to minors.

Throughout the experiment, you will represent the city councils of various cities, each with different preferences concerning smoking, taxation, campaigns, and financial possibilities available to implement different solutions. You will negotiate with another professional mediator, who will represent local business owners of the same cities, which may also differ in their preferences. In each negotiation, you should try to negotiate a solution that is as close as possible to the preferences of the city that you represent. To help you in this task, you will receive a "preference card" for each city. This card lists the different options for each of the four issues. The background color of each option shows whether the city council views these options as positive or negative according to the gradient bar below.

The options that are in line with the goals of the council (and that you should try to negotiate for) are indicated with a blue background. An orange background, on the other hand, shows that the option is contrary to the goals of the city council. These options would therefore make a proposal less appealing for the city you represent. Finally, options with a white background are neutral options that the council considers to be neither positive nor negative. An agreement with only orange options will be unacceptable for the city council you represent. However, orange options can be compensated by blue options. For example, a bright blue option on taxation combined with a slightly orange option on the anti-smoking campaign results in an acceptable agreement.

Your aim is to negotiate a mutual agreement. If you believe it is not possible to negotiate an agreement with your trading partner, you may decide to stop the negotiation and move on to the next. During the negotiation, it is important to remember that the preference cards you receive are private information. As a professional mediator, you are not allowed to show this information to your trading partner.

8.2 Business owners representative cover story

In this experiment, you will play the role of a professional mediator employed by the an organization for small business owners. In this capacity, your task is to negotiate the implementation of new government regulations concerning an anti-smoking campaign. Although the campaign is nation-wide, cities are allowed to make their own decisions concerning the details of the implementation. In

fact, different cities turn out to be quite varied in their preferences. For this reason, you will enter separate negotiations with your trading partner for each city. During these negotiations, the following four aspects need to be decided on:

- Scope of smoking ban;
- Taxation of tobacco products;
- Anti-smoking campaign; and
- Enforcement of tobacco sales restrictions to minors.

Throughout the experiment, you will represent the small business owners of various cities, each with different preferences concerning the sales of tobacco, taxation, advertising, and government control in their stores. You will negotiate with another professional mediator, who will represent the city councils of the same cities, which may differ in their preferences. In each negotiation, you should try to negotiate a solution that is as close as possible to the preferences of the business owners that you represent. To help you in this task, you will receive a ``preference card'' for each city. This card lists the different options for each of the four issues. The background color of each option shows whether the business owners view these options as positive or negative according to the gradient bar below.

The options that are in line with the goals of the business owners (and that you should try to negotiate for) are indicated with a blue background. An orange background, on the other hand, shows that the option is contrary to the goals of the business owners. These options would therefore make a proposal less appealing for the business owners you represent. Finally, options with a white background are neutral options that the business owners considers to be neither positive nor negative. An agreement with only orange options will be unacceptable for the organization you represent. However, orange options can be compensated by blue options. For example, a bright blue option on taxation combined with a slightly orange option on the enforcement of tobacco sales results in an acceptable agreement.

Your aim is to negotiate a mutual agreement. If you believe it is not possible to negotiate an agreement with your trading partner, you may decide to stop the negotiation and move on to the next. During the negotiation, it is important to remember that the preference cards you receive are private information. As a professional mediator, you are not allowed to show this information to your trading partner.

8.3 Scenarios

Our setting includes nine different settings that differ in their mix of cooperative and competitive elements. The preferences of the representatives of the small business owners and the city council are listed in Table 4. Scenarios 1, 5, and 7 are cooperative scenarios, in which the preferences of both negotiators mostly align. Scenarios 3, 6, and 9, on the other hand, are competitive scenarios, in which the preferences of both negotiators are mostly contrary to one another. The remaining scenarios 2, 4, and 8 are more neutral in the sense that preferences sometimes align.

	1. Scope				2. Tax					3. Campaign				4. Enforcement				
	a	b	c	d	a	b	c	d	e	a	b	c	d	a	b	c	d	e
Scenario 1																		
Council	-4	-1	2	1	-1	1	2	0	-2	-2	2	-1	-3	2	0	-1	-2	-4
Business	0	1	2	-1	4	3	0	-1	-2	-4	2	0	0	1	-2	-4	-4	-4
Scenario 2																		
Council	-1	0	2	4	0	-1	-2	-3	-4	0	-1	1	3	0	1	-1	2	4
Business	2	1	0	-1	2	1	0	-3	-4	-1	2	3	-3	0	1	-2	-3	-1
Scenario 3																		
Council	-2	2	3	1	-4	-3	0	3	1	3	4	0	-2	-4	-2	0	-1	2
Business	4	1	-1	-4	4	-1	-3	-4	-4	0	1	4	2	4	-1	1	-2	-4
Scenario 4																		
Council	2	-1	-2	-4	-4	-2	0	2	4	0	0	2	1	0	-2	0	-3	-4
Business	4	2	1	0	3	0	-2	-3	-4	0	-1	1	0	1	0	-1	-4	1
Scenario 5																		
Council	-4	2	1	0	-2	4	2	-2	-4	0	0	2	1	0	-2	0	-3	-4
Business	2	0	0	-4	4	3	0	-2	-4	-2	-4	1	0	2	-3	0	-2	-4
Scenario 6																		
Council	-4	2	1	0	-4	2	4	2	-4	3	1	-2	-4	-2	1	-2	2	-2
Business	2	1	1	-1	4	-1	-2	-4	-4	-3	-1	4	2	2	1	-2	0	2
Scenario 7																		
Council	-4	-1	3	0	4	0	-1	-2	-4	0	2	0	-4	2	0	2	0	-4
Business	-2	0	0	-4	0	0	-2	-4	-4	0	0	4	3	4	0	-2	-2	-4
Scenario 8																		
Council	-4	-3	-2	0	2	1	0	-2	-4	-2	0	2	1	0	-1	0	0	2
Business	0	2	1	-2	2	1	0	-1	-4	-1	1	-4	-4	0	2	-2	3	-1
Scenario 9																		
Council	-4	0	1	3	0	1	2	3	4	0	0	-2	-4	0	4	-2	-1	-4
Business	4	1	-1	-4	4	0	-1	-2	-3	-2	-3	0	0	2	0	-2	-3	4

Table 4: Preferences for both negotiating parties in each of the nine scenarios

8.4 Results of the pilot experiment

In this section, the results of the pilot experiment are listed in full. Audio recordings and tran- scripts of the experiments are available from the Metalogue ownCloud server.

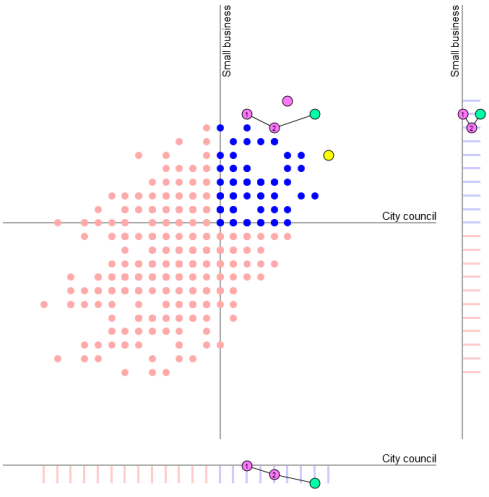

(a) Scenario 1

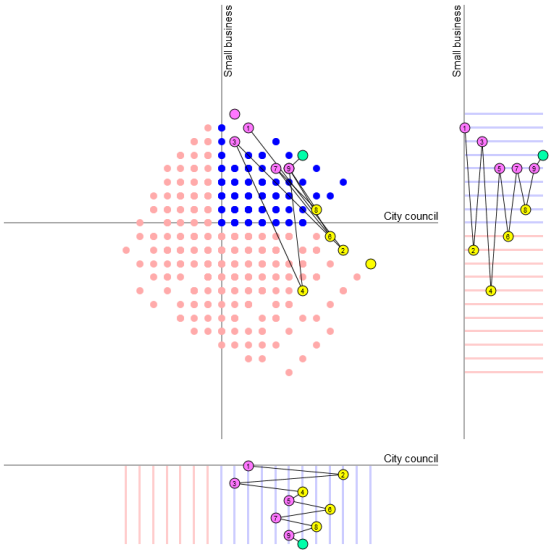

(b) Scenario 2

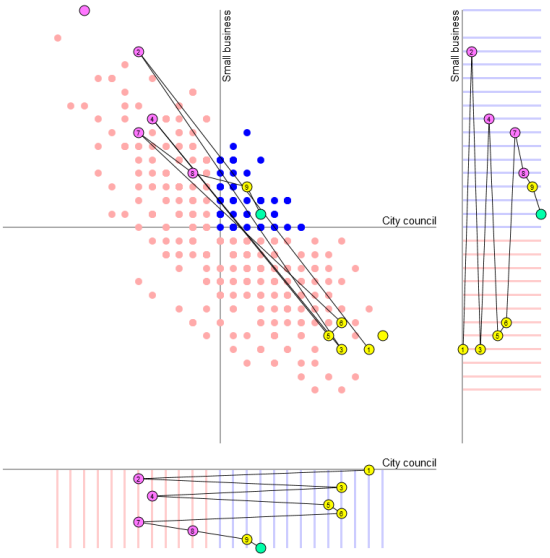

(c) Scenario 3

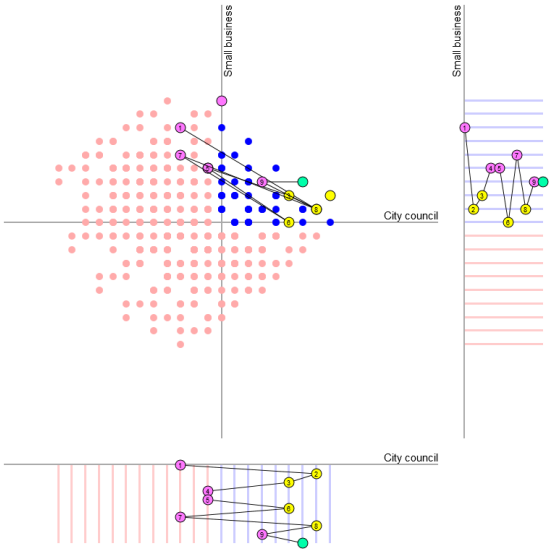

(d) Scenario 4

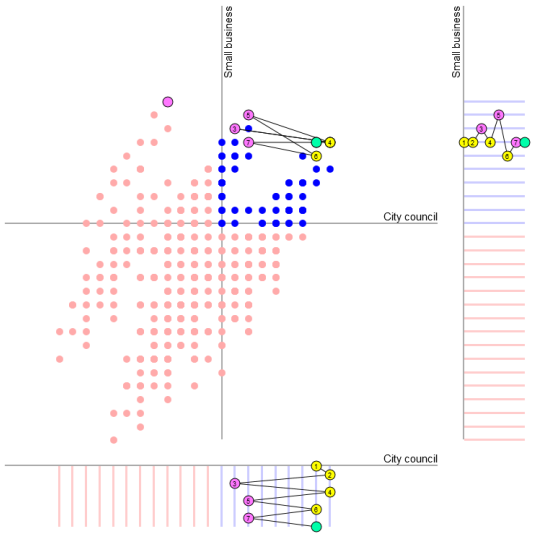

(e) Scenario 5

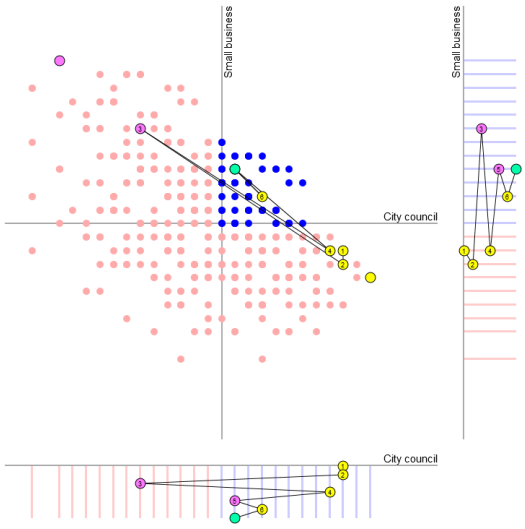

(f) Scenario 6

Figure 3: Negotiation space and negotiation outcome for scenarios 1–6. The top left panel in each figure shows the possible negotiation outcomes in terms of the score for the small business owners representative (vertical) and the city council representative (horizontal). Blue dots indicate outcomes that are acceptable to both negotiator, while red dots indicate unacceptable outcomes. Large, unnumbered dots show the best possible outcomes for the small business owners representative (purple) and the city council representative (yellow), as well as the final agreement (teal). The sequence of offers made during the negotiation is indicated by connected, numbered dots. The color of these dots shows which negotiator made the offer. The panels on the right and the bottom show the same information from the perspective of the small business owners representative and the city council representative, respectively.

© METALOGUE Consortium [2015]