

Using Cognitive Modeling to Construct a Believable Opponent for the game of 'No Thanks!'

Bachelor Thesis

Wolter Peterson – University of Groningen
w.peterson@student.rug.nl

In a world in which computer games play an ever growing role there is an increasing demand for interesting artificial intelligence (AI) opponents in games. To date many games provide AI players. However most of the time these AI's fail in creating an interesting gaming experience for the human player. One of the reasons for this could be that computer players use techniques so different from those of humans they seem very unrealistic and resemble cheating. To test this hypothesis we modeled several opponents for the card game 'No Thanks!' based on strategies reported by human players. These opponents differed in how much they take into account the entire game and all players. A player that does this more strongly indeed performs better in the game. However, having a model that is able to play better is not a guarantee for it being interesting as an opponent.

In most games a player competes against several opponents. In some cases, one or more of these opponents are computer opponents. When playing these games people tend to want a challenging opponent. An opponent that is fun to play against. This brings us to the question: What is a challenging opponent? The classical view about this is that an opponent should make good moves.

The most well-known example of playing a game, which does not have its origin as a computer game, against a computer opponent is chess. Chess programs make well calculated, good moves, and they play according to the rules. This should translate to these programs being challenging opponents. However, a lot of people do not find playing against such an opponent enjoyable. Chess programs are hard to beat and often employ brute-force techniques, allowing them to calculate a large amount of moves and outcomes. The only way to beat such an opponent is by finding a weakness in the program and exploiting this.

Apparently making good moves is not enough to be a challenging and fun opponent. People want opponents that can beat them one time but lose the next, while still making realistic moves. Moves they could have done themselves.

Taatgen, van Oploo, Braaksmā and Niemantsverdriet, (2003) used the game of SET to test the notion that a program that mimics human playing is enjoyable to play against. Based on experiences from players and a short experiment, they implemented a cognitive model. By letting participants play against their model and an opponent that waited a set time before it found a set, they showed that people indeed find it more enjoyable to play against an opponent using a human playing style.

SET is a discrete game in which each set

of cards can be seen as a separate round. However, most games are continuous, and current actions are influenced by earlier events. Previous research (e.g. Lebiere & West, 1999) showed that in such situations it is important to be able to predict what your opponent might do. As if this is not all hard enough, most games are imperfect information games. In imperfect information games it is, in principle, not possible to calculate the optimal move as not all information is available to the player.

A typical machine approach to playing games would mean combining a lot of calculations, for games like chess, with true randomness, for imperfect information games. Humans however, seem to play quite differently. West, Lebiere and Bothell (2006) argued that there is a difference between an optimal playing style and a maximizing playing style. Optimal here means optimal as it is meant in classic game theory. In game theory an optimal strategy is a strategy in which (1) moves can be described as being selected based according to fixed probabilities and (2) these probabilities describe an optimal or near-optimal approach to the game. Maximizing is used to refer to a strategy that uses learning, reasoning and problem solving to respond in adaptive ways, and trying to improve the choice of future moves. West, et al (2006) argued that humans use a maximizing playing style. Human players tend to look for patterns, real or only perceived, and try to exploit these patterns, instead of calculating which move will have the optimal outcome or maximum gain.

For this study we will combine the methods employed by Taatgen et al. (2003) with knowledge about human game playing. We will attempt to create a cognitive model able to play the game of 'No Thanks!' in an interesting and

believable way, and we try to test the notion that a good player indeed has to reason about opponents and be able to predict what they might do. To this end, two models will be implemented using ACT-R (Anderson & Lebiere, 1998). The first model will only look at its own situation, while the other model takes into account the entire game and the states of the other players.

Both the term "interesting" and the term "believable" are vague terms which are hard to specify. For this study the term "believable" was linked to how humanlike the models played. We did not specify the term "interesting", but instead kept this term subjective and open to interpretation by the participants.

The Game of 'No Thanks!'

'No Thanks!'¹ is a sequence collecting game for three, four or five players. The game is played with a deck of cards consisting of cards with values between 3 and 35 (one of each), from which nine cards are removed at random at the start of each game. Players start the game with 11 chips in their hand. These chips can be used to do actions with, and have to be kept so that the other players can't see how many there are left.

During a round one card from the deck is revealed. The player who currently is the starting player has to decide whether to take the card or place one of his chips on it to pass. This continues until one of the players decides to take the card. This player then takes the card and all the chips on it, placing the card in front of him on the table. After this a new card is revealed. This goes on until the deck of cards is empty. When the deck is empty, the scores are calculated. The values of the cards each player took are added up, and from this the amount of chips each player has left is deducted. The goal of the game is to end with as few points as possible. When a player has a sequence of cards, points are only gained for the lowest card in this sequence. If, for example, a player has cards 10 and 11, only the 10 is counted towards his score. Sequences can be of any length.

Obtaining Strategies

We hypothesize that an opponent that is interesting to play against is an opponent that mimics human playing. To be able to see whether this is true or not, a model that plays in a humanlike manner was needed. For this it was necessary to know how humans play the game.

Based on own prior knowledge and initial observations a questionnaire was prepared to enquire about aspects and strategies in the game.

The Questionnaire

The questionnaire was roughly divided into five parts. The first part was a single question about whether or not players thought the number of players influenced the game and, if so, in which manner. This question showed the number of players had a large influence on the game. More players means more competition for the cards, and a different amount of chips being played and desired. Because of this, and based on the recommendations of players, the number of players was fixed to four for the rest of this study.

Next were questions about card preferences. We asked participants which cards they tended to take in the beginning (No preference was reported), whether they would take a card that was beneficial to another player but that would hurt themselves (Most players would not), and what other factors might play a role (The number of chips left was mentioned in a lot of the cases).

The third group of questions was about the amount of chips players preferred and in what way players kept track of the chips. Most players keep track of the, exact, amount of chips on the card. Only a few players reported to keep their knowledge about this to an estimate of few or many. The chips the other players have are harder to keep track of, and because of this most people just make an estimate (Few, Many, Nearly out, out), while both not keeping track at all and keeping track of the exact amount got picked only once. Players also reported that the number of chips they had left influenced how likely they were to take a certain card. Less chips often meant players took a card with fewer chips on it, just to not run out of chips entirely. We also asked them how greedy they were. In other words, how much would they pass on a card that was already beneficial, in order to get more chips. One group reported not really doing this, while others did. How much they would pass depended on the card and on how the other players played. If the other players would take cards quickly, people tended to be less greedy. Finally we asked them to list how many chips they would want on a certain card. The values they were asked to give their desired amount of chips for were; 5, 10, 15, 20, 25, 30. The results from this question can be found in Table 1 and will be discussed further in the next section.

¹ No Thanks! is a game by Z-man games (www.zmangames.com)

The fourth group of questions were open questions about feedback and how players interact with and react to other players. We asked how they decided whether the decision they made was a good one. This turned out to be hard to answer and most players did not do this, at least not consciously. Players did view passing on a card to get more chips a bad move when another player took the card before it got back to them. In line with that it is hard to keep track of someone else his amount of chips, quite a few, mainly novice, players reported they did not really take into account what other players might do for a certain card. They reported just playing their own game, and made decisions based just on their own cards and chips.

The last part of the questionnaire contained two multiple choice questions about expertise, namely: “How often do you play the game”, and “How would you categorize your level of experience”. These questions revealed that our initial participants were all novice players with little or no experience in the game. Because of this no general conclusions could be made yet. To circumvent this problem the questionnaire was distributed on a boardgame community website, www.boardgamegeek.com, by posting it on the forums and contacting several members there who have logged a lot of plays in the game. Combining the answers of the novice and experienced players several main strategies were uncovered.

Card	Calculated 1/3	Reported Average / 3
5	1.67	1.6
10	3.33	3.6
15	5	4.6
20	6.67	6.4
25	8.33	8.14
30	10	11.43

Table 1: Summary of reported amount of chips players want on middle cards with a certain value

The Strategies

The answers to the questionnaire revealed two major strategies, or possible generalizations. The First interesting generalization that was possible was about the amount of chips people preferred. Taking into account some slight differentiations, people tended to think of taking a card, that was not part of a sequence of cards they already had, as “not too bad” approximately when the number of chips on the card was 1/3 of the value of the card (See Table 1). The deviation from this 1/3-rule was influenced by the amount of chips players had left and how their opponents

played. Willingness to take cards at these numbers of chips implies players are willing to incur a bigger loss for taking a high card than for taking a low card.

The second, and perhaps most interesting difference that was observable was the difference in in how far a player reasoned about its opponents. Some players only take into account their own situation, which in the case of this game means their own cards and number of chips. Another group of players did not just look at themselves, but also at the cards of their opponents. These players also tried to guess how many chips the other players were likely to have. The players who only looked at their own situation tended to pick up a card as soon as it was beneficial, or not too bad. The other group tried to maximize the amount of chips they gained by reasoning whether or not their opponents would take the card or whether it would get back to them once more, having 4 more chips on it.

The Model

Although the game only has two possible actions a player can take, several different situations can occur. Each elicits a different responds and requires a different approach. For the implementation of our models we divided the possible situations that can occur with respect to the card currently up for grabs into five categories; filling a gap, being of value one lower than currently owned cards, being of value one higher than currently owned cards, being of value two higher or lower than already owned cards, and being further removed in value from already owned cards. These situations differ in how much the card that can be taken is beneficial for the player to take.

In the category where a gap can be filled, the card in the middle has a value that lies between two cards the player already has. If, for example, the card in the middle is card 23 and the player already has 22 and 24, the situation would fall in this category. In this case it would be very beneficial for the player to take the card even if there are no chips on it. The score gained from taking this card is equal to the value of the higher of the two cards the player already owns.

In the category where the card in the middle is of value one lower, it is also beneficial to take the card even when there are no chips on it. In this case a one point gain is achieved. When the card in the middle is of value one higher than an already owned card, no loss is incurred for taking it. In this scenario it is beneficial to grab

the card as soon as there are one or more chips on it.

In the last two categories it is less beneficial, at least in the short term, to take the card. When taking a card that has a value of two higher or lower, a gap is left to fill. Whether or not it was a good choice to take this card therefore depends very strongly on which cards show up later on. When the card in the middle has a value that is not near the value of any card the player has, it is least beneficial to take the card. The chance to turn this card into a sequence with the cards the player already has is very low. Taking such a card, specially later in the game, often means incurring a loss equal to the value of the card.

Based on the strategies we obtained and the categories described above, two models were implemented. The big difference between the models lies in the in how much they reason about their opponents. One model only looks at its own situation while the other takes into account the other players as well.

The Basic Model

The first model, which we call the basic model, only looks at its own situation. This means a player controlled by this model will make its decision to take a card or pass on it based on which cards and how many chips it has itself.

When this model encounters a situation that falls in one of the first two categories, filling a gap of being one lower than already owned cards, it will immediately take the card. It will improve its own score by doing this. If the card in the middle is of value one higher than a card this model already has, it will take the card when there are one or more chips on there. For the other two categories it will decide how many chips it wants, mainly based on the 1/3 rule, and will take the card as soon as the number of chips on the card are equal to or higher than the number of chips wanted.

The Reasoning Model

The second model that was implemented is the reasoning model. As opposed to the basic model, this model will not instantly take a card as soon as its desired number of chips is reached. Instead, in this case this model will try to determine what the other players will do. If the model thinks the other players' desired amount of chips will not yet be reached by the time it's their turn, it will decide its save to pass. If this "thinking about others" goes well the other

players will indeed pass and the card will get back to the player, thus increasing the amount of chips gained and therefor decreasing the loss incurred for taking the card.

For example, the player has card 23 and number 22 is turned open in the middle. The basic player would just take this card, bettering its own score by one point. The reasoning player will, if none of the other players has card 21, reason that it is save to pass on this card for now. This means it will have four chips when it is this players' turn again, leading to a gain of three chips and therefore a four point increase in score. The player might then decide to take the card of pass on it again, perhaps even further bettering the score.

Details on the Models

For the implementation of the models a java implementation of ACT-R, created by Dario Salvucci, was used. This implementation combined the editor, output and a task interface panel into the same window of the program. With this java implementation it is not necessary, or even possible, to declare chunk types. The following aspects of ACT-R are relevant for this model:

- The model uses the goal buffer, the retrieval buffer (the buffer where chunks recalled from the declarative memory are stored), and the imaginal buffer (which can be seen as a sort of working memory) to store relevant information during the task.
- All actions and operations of the model are controlled by production rules. Based on the contends of the goal buffer, the retrieval buffer and the imaginal buffer the appropriate production is chosen and its actions are executed. When multiple models play against each other, each has its own productions for the specific reasoning they do while they share common productions like clicking the buttons.
- The declarative memory is used to store the decisions the players made. These decisions can then later be recalled to know which action to take.
- The motor module is used to take the actions possible in the game, i.e. clicking the take or pass buttons. Although this is not necessary for the functioning of the model it does provide a more engaging experience for the human player. Using the motor module means a cursor will move according to the actions the model takes.

- To enable the motor module to click the buttons the visual module is used. Other information that is necessary such as which card is in the middle, how many chips are on the card in the middle, which players' turn it is and what cards the players have, is requested directly from the task. Although this could also be read with the visual module, the process of looking and the time this requires are not important for this task.
- Information which requires calculations, such as determining $1/3$ of the value of the card, is also obtained directly from the task. This approach was chosen because the goal was to create a model that can play the game, and not a model that can do math.

When not looking at the details of the reasoning, a turn of a model player is played as following:

1. The player will look whether or not it has more than zero chips left. If this is not the case the model will take the card. If the model has more than zero chips left, then
2. The model will look at the card in the middle and which cards it has itself. Based on this the model decides how many chips it wants on the card. If the model is of the reasoning kind it will then do the same for the other players.
3. Based on the amount of chips it wants, and possibly the amount the other players want the model decides to either pass or take the card.
4. When the decision is made to either take or pass the card, the model clicks the appropriate button. After this the task handles the updating of the game, and then it's the next players' turn.
5. The decision is saved in the declarative memory. If the card in the middle has not changed on the next turn of this player, the previously made decision is recalled and used again to determine whether or not to take or pass.
6. To determine whether to take or pass the card, the model will check whether or not the amount of chips on the cards meets its desired amount. If this is not the case, the model passes. If the desired amount is met the basic model will take the card. The reasoning model will check its desired amount against the desired amount of the other players. If it is save to pass it will

pass. Otherwise it will take the card.

Model versus Model

One of the questions we try to answer is whether or not a player that takes into account the entire game and reasons about its opponents will do better than a model that only looks at itself. We hypothesize that a model that does this reasoning will perform better.

In order to test this the game was set up in such a way that two players of the reasoning type and two players of the basic type played against each other. The models played 500 games and for each of these 500 games we logged the scores of each player, as well as who won. Throughout these games the order in which the models were situated were kept the same. i.e. each player had a fixed playing-style.




Figure 1: Total amount of wins per player in 500 games of 'No Thanks!'

Results

Figure 1 shows the total amount of wins per player. In this figure it becomes apparent that the reasoning model wins more games than the basic model. Figure 2 shows the amount of wins after combining the players into which model they represent, and dividing the 500 games into 10 blocks of 50 games. On average the reasoning model won 31,10 out of every 50 games, against just 19,20 wins for the basic model. The reason that this number is slightly over 50 is because of ties between the models in some of the games.




Figure 2: Wins per model, per block of 50 games

Not just wins are important in this game. Although a lower score most likely means winning, it could be possible that both models score almost the same, but that the reasoning model just gets “lucky” most of the time. Figure 3 shows the mean score of each model per 50 game block. Here we find that the basic model has a mean score of 58.55 (SD = 58.55) over all the games, against a score of 49.67 (SD = 22.837) for the reasoning model. That is a difference of 15%.




Figure 3: Average score per model, per block of 50 games

Finding both a large difference in wins and in scores, we can conclude the reasoning model does better in this game than the basic model.

Model versus Human

The second, and perhaps most important, question we try to answer is whether or not we managed to construct an interesting and believable opponent for this game.

Participants

In total 11 participants, 5 female and 6

male, were used in the Model versus Human comparison. 7 were artificial intelligence students, 2 were college students and 2 were working adults. Ages ranged from 18 to 23, with two outliers of 50+. The mean with the outliers is 27, and without it 21.2. All participants were novice players, and except for three they were all new to the game.

Procedure

For the experiment we let the participants play against three model opponents. Of these opponents one was a basic model and two were reasoning models. Each participant played two to four games, on average three, against these opponents. After this they were asked to rate, on a scale of 1-5, how interesting and how humanlike they thought the opponents were. These ratings are independent of each other (i.e. multiple players can have the same rating). The participants had no information about the opponents other than that which they observed during playing the game. Of the human versus model games we also logged the scores and who won. The position of each player was fixed and kept the same throughout all games.

Wins and Scores

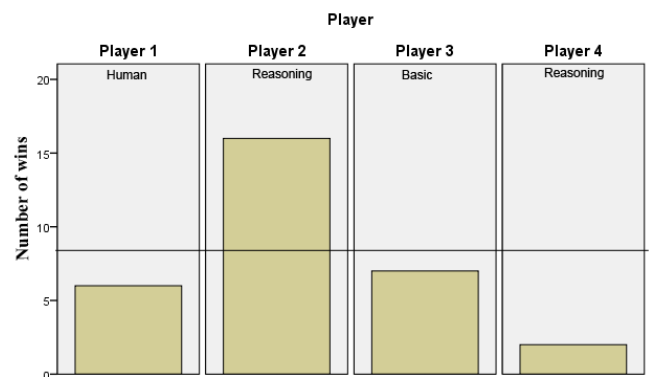


Figure 4: Total amount of wins per player for 31 games

Figure 4 shows the total number of wins per player for the 31 games that were played. Based on what we saw in the Model versus Model comparison we’d expect the two reasoning players to do better than the basic model. Figure 4 shows that this is not entirely the case. Although player 2 wins a lot, and combined player 2 and 4, the reasoning models, still win more than the chance level of 8 wins (32 games divided by four players). Player 4 does a lot worse than both the human player and the basic model. We also see that the human players win less than the basic model. This distribution did not occur by chance and is therefore the result of the way the players played the game, $\chi^2(3, N=124) = 18,02, p = 0$.




Figure 5: Boxplot of the distributions of the scores per player.

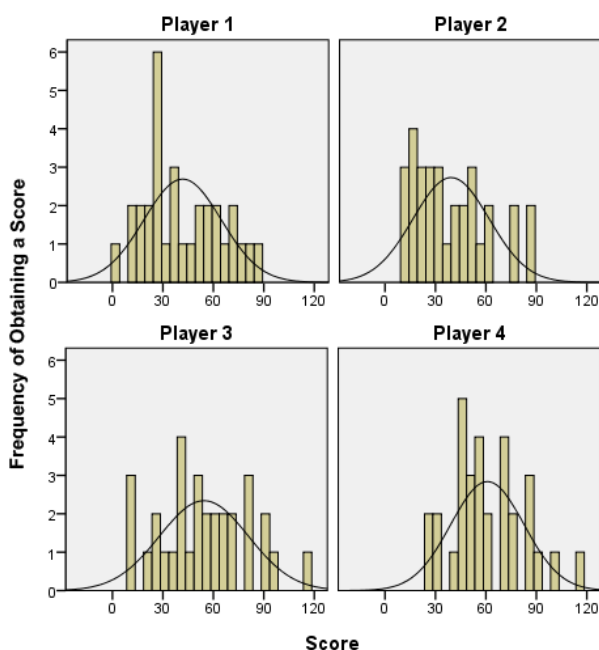


Figure 6: Distributions of the score per player. Horizontal axis contain score per 5 point interval. The curve is the distribution curve.

Looking at the scores (Figure 5 & 6) we see a slightly different picture. Again player 2 does best (Mdn = 33) while player 4 comes in last (Mdn = 57). Here however, the human player (Mdn = 36) seems to do better on average than the basic model, player 3 (Mdn = 53), and almost as good as player 2. The data for player 1 seems to be normally distributed. The distribution is slightly skewed to the lower scores, with $M = 40.94$, $SD = 22.99$, because of the large spike of scores between 25 and 30. This might indicate the human player scores randomly over all the games. Player 2 is skewed more strongly to the right ($M = 38.25$, $SD = 22.65$). Attaining a low score in most games. Figure 5 shows player 2 has a long upper quartile with a shorter lower quartile, and the sample maximum also is further away from the

upper quartile than the sample minimum is from the lower quartile. The real mean score is therefore probably lower than the current reported mean and median. This difference in scores, coupled with bad luck for the participants might have led to the big difference in wins.

The Ratings

If the ratings of how interesting and humanlike were largely influenced by the type of model behind a player, player 2 and player 4 should show similar ratings.

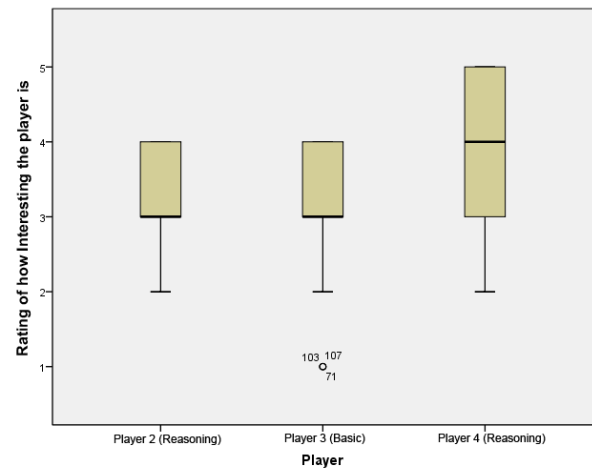


Figure 7: Per player boxplot of how interesting the opponent was.

Figure 7 shows the ratings of how interesting each of the players was found by participant. Here we see that although player 4 was the worst player in terms of scores and wins, it is rated slightly more interesting than the other two players. Using a Kruskal-Wallis Test no significant difference between the ratings was found, $\chi^2(2, N=33) = 2.30$, $p = 0.32$.

There is only a minimal difference between player 2 and player 3. The only difference stemming from a few very low ratings for player three.




Figure 8: Boxplot of how humanlike the opponents were.

Figure 8 shows the ratings of how humanlike participants thought the playing styles used by their opponents were. This figure shows the same effect as figure 7, with the worst scoring player being rated most humanlike. Again the ratings for player 2 and player 3 are very similar. The differences in rating for how humanlike the players were found are also not significant, $\chi^2(2, N=33) = 2.85, p = 0.24$.

Conversations with the participants after they had played the games and filled in the ratings, revealed that participants perceived behavior that was not there. The way in which cards came up and were taken by the models led to participants to think their opponents did certain things for a different reason than they actually did. A lot of participants reported player 4 taking cards just to mess with them, while such a thing is not part of either of the models. This indicates the ratings might be influenced by the outcomes of the games.

Discussion and Conclusion

For this project we posed two questions. (1) can we create an interesting and believable opponent for the game of 'No Thanks!', and (2) is an player that takes into account the entire game and which reasons about the other players, better than a player who only looks at its own state.

In the Model versus Model section we saw that the reasoning model scores significantly better and wins significantly more than the basic model. Based on this we can conclude that indeed a player that takes into account the entire game situation is better than a player that only looks at itself. For this it is needed that the player has a clear understanding of its opponents. How this understanding can be gained, however, goes beyond the scope of this project.

The first question, unfortunately, cannot be answered in full. The ratings obtained from participants are close together. Conversations revealed participants observed behavior that was not part of the models, rating the opponents for the different reasons than we'd expect. These findings perhaps hint towards humans not only failing to use optimal strategies, but that they actually find it non-humanlike when their opponent does manage to do so.

Further inquiry with people who played the game against the models more often revealed they manage to find the behavior behind the models. This same effect occurred when participants played more than just a few games

against the same human opponents. There seems to be an effect of meta-gaming going on. People tend to look for patterns over the course of several games. This effect is most likely closely linked to the question of how a model can learn to predict its opponents.

The position in which a player is situated also seems to play a role. Based on what was observed in the model versus model condition, the reasoning players should win approximately the same amount of games. However, when a different player, the human, is introduced one of the reasoning models does a lot worse, while the other scores equally good or even better than before. Why this is remains unclear.

References

- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.
- Lebiere, C., & West, R. L. (1999). A dynamic ACT-R model of simple games. *In Proceedings of the Twenty-first Conference of the Cognitive Science Society* (pp. 296-301.) Mahwah, NJ: Erlbaum
- Taatgen, N.A., van Oploo, M., Braaksma, J. & Niemantsverdriet, J. (2003). How to construct a believable opponent using cognitive modeling in the game of set. In F. Detje, D Dörner and H. Schaub (eds.), *Proceedings of the fifth international conference on cognitive modeling* (pp. 201-206). Bamberg: Universitätsverlag Bamberg.
- West, R. L. , Lebiere, C. & Bothell, D. J. (2006). Cognitive architectures, game playing, and human evolution. In R. Sun (ed.), *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*. Cambridge University Press; New York, NY, 103-123.