Two Level Recursive Reasoning by Humans Playing Sequential Fixed-Sum Games

Adam Goodie
Dept. of Psychology
University of Georgia
Athens, GA 30602
goodie@uga.edu

Prashant Doshi
Dept. of Computer Science
University of Georgia
Athens, GA 30602
pdoshi@cs.uga.edu

Diana Young
Dept. of Psychology
University of Georgia
Athens, GA 30602
dlyoung@uga.edu

ABSTRACT

Recursive reasoning of the form *what do I think that you think that I think* (and so on) arises often while acting rationally in multiagent settings. Previous investigations indicate that humans do not tend to ascribe recursive thinking to others. Several multiagent decision-making frameworks such as RMM, I-POMDP and the theory of mind model recursive reasoning as integral to an agent's rational choice. Real-world application settings for multiagent decision making tend to be *mixed* involving humans and human-controlled agents. We investigate recursive reasoning exhibited by humans during strategic decision making. In a large experiment involving 162 participants, we studied the level of recursive reasoning generally displayed by humans while playing a sequential fixed-sum, two-player game. Our results show that subjects experiencing a strategic game made more competitive with fixed-sum payoffs and tangible incentives predominantly attributed first-level recursive thinking to opponents. They acted using second level of reasoning exceeding levels of reasoning observed previously.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms

Theory, Performance

Keywords

decision making, recursive reasoning, human behavior

1. INTRODUCTION

Strategic recursive reasoning of the form *what do I think that you think that I think* (and so on) arises naturally in multiagent settings. For example, a robotic uninhabited aerial vehicle (UAV)'s decision may differ if it believes that its reconnaissance target believes that it is not being spied upon in comparison to when the UAV believes that its target believes that it is under surveillance. Specifically, an agent's rational action in a two-agent game often depends on the action of the other agent, which, if the other is also rational, depends on the action of the subject agent.

Assumptions of *common knowledge* [11, 10] of elements of the game tend to preclude the emergence of recursive reasoning. However, not all elements can be made common knowledge. For example, an agent's belief is private especially in a non-cooperative setting. Multiple decision-making frameworks such as the recursive modeling method (RMM) [17, 18] and interactive partially observable Markov decision process (I-POMDP) [16, 8] model recursive beliefs as an integral aspect of agents' decision making in multiagent settings.

Real-world applications of decision making often involve mixed settings that are populated by humans and human-controlled agents. Examples of such applications include UAV reconnaissance in an urban operating theater and online negotiations involving humans. The optimality of an agent's decisions as prescribed by frameworks such as RMM and I-POMDP in these settings depends on how accurately the agent models the strategic reasoning of others. A key aspect of this modeling is the depth of the recursive reasoning that is displayed by human agents.

Initial investigations into ascertaining the depth of strategic reasoning of humans by Stahl and Wilson [25] and more recently, by Hedden and Zhang [21] and Ficici and Pfeffer [12] show that humans generally operate at only first or second level of recursive reasoning. Typically the first level, which attributes no recursive reasoning to others, is more prominent. Evidence of these shallow levels of reasoning is not surprising, as humans are limited by bounded rationality.

Increasing evidence in cognitive psychology [14, 15, 19] suggests that tasks which are ecologically more valid and which incorporate tangible incentives induce decisions in humans that are closer to being rational. Both these aspects were lacking in the experiments conducted by Hedden and Zhang [21]. We hypothesize that a strategic setting that is realistic, competitive and includes tangible incentives would increase participants' tendency to attribute levels of reasoning to others that reflect individuals' actual level of reasoning.

In this paper, we report on a large study that we conducted with human subjects to test our hypothesis. We constructed a task that resembled the two-player sequential game as used by Hedden and Zhang but made more competitive by incorporating fixed-sum outcomes and monetary incentives. Subjects played the game against a computer opponent, although they were led to believe that the opponent was human. Different groups of subjects were paired against an opponent that used no recursive reasoning (zero level) and opposite one that used first-level reasoning. We also manipulated the realism of the task between participants, with one group experiencing the task described abstractly while the other experienced a task that was structurally identical but described using a realistic cover story involving UAV reconnaissance.

Data collected on the decisions of the participants indicate that (i) subjects acted accurately significantly more times when the opponent displayed first-level reasoning than when the opponent was

AAMAS 2009 Workshop on Multi-agent Sequential Decision-Making in Uncertain Domains, May 11, 2009, Budapest, Hungary.

at zero level. The particpants also learned the reasoning level of opponents more quickly than reported previously by Hedden and Zhang. (ii) No significant difference in the accuracy of the decisions was noticed between abstract and realistic task settings. Thus, our study reveals clear evidence that higher levels of recursive reasoning could be observed in humans under simpler and more competitive settings with tangible incentives. However, there is room for future research on whether increased realism contributes to deeper levels of strategic reasoning.

2. RELATED WORK

Harsanyi [20] recognized that indefinite recursive thinking arises naturally among rational players, which leads to difficulty in modeling it computationally. In order to, in part, avoid dealing with recursive reasoning, Harsanyi proposed the notion of agent types and common knowledge of the joint belief over the player types. However, as shown in [11, 10], common knowledge is itself modeled using an indefinite recursive system.

Since Harsanyi's introduction of abstract agent types, researchers have sought to mathematically define the type system. Beginning with Mertens and Zamir [24], who showed that a type could be defined as a hierarchical belief system with strong assumptions on the underlying probability space, subsequent work [5, 22] has gradually relaxed the assumptions required on the state space while simultaneously preserving the desired properties of the hierarchical belief systems. Along a similar vein, Aumann defined recursive beliefs using both a formal grammar [1] and probabilities [2] in an effort to formalize interactive epistemology.

Within the context of behavioral game theory [6], Stahl and Wilson [25] investigated the level of recursive thinking exhibited by humans. Stahl and Wilson found that only 2 out of 48 (4%) of their subjects attributed recursive reasoning to their opponents while playing 12 symmetric 3×3 matrix games. On the other hand, 34% of the subjects ascribed zero-level reasoning to others. Remaining subjects utilized either Nash equilibrium based or dominant strategies. Corroborating this evidence, Hedden and Zhang [21] in a study involving 70 subjects, found that subjects predominantly began with first-level reasoning. When pitted against first-level co-players, some began to gradually use second-level reasoning, although the percentage of such players remained generally low. Hedden and Zhang utilized a sequential, two-player, general-sum game, sometimes also called the Centipede game in the literature [4]. Ficici and Pfeffer [12] investigated whether human subjects displayed sophisticated strategic reasoning while playing 3-player, one-shot negotiation games. Although their subjects reasoned about others while negotiating, there was insufficient evidence to distinguish whether their level two models better fit the observed data than level one models.

Evidence of recursive reasoning in humans and investigations into the level of such reasoning is relevant to multiagent decision making in *mixed* settings. In particular, these results are directly applicable to computational frameworks such as RMM [17], I-POMDP [16] and cognitive ones such as theory of mind [9] that ascribe intentional models of behavior to other agents.

3. EXPERIMENTAL STUDY: HIGHER LEVEL RECURSIVE REASONING

In a large study involving human subjects we investigate whether subjects would generally exhibit a higher level of recursive reasoning under particular settings that are more typical of realistic applications.

We begin with a description of the problem setup followed by the participating population and our methodology for the experiment.

3.1 Problem Setting

In keeping with the tradition of experimental game research [7, 6] and the games used by Hedden and Zhang [21], we selected a two-player alternating-move game of complete and perfect information. In this sequential game, whose game tree is depicted in Fig. 1(a), player *I* (the leader) may elect to *move* or *stay*. If player *I* elects to move, player *II* (the follower) faces the choice of moving or staying, as well. An action of stay by any player terminates the game. Note that actions of all players are perfectly observable to each other. While the game may be extended to any number of moves, we terminate the game after two moves of player *I*.

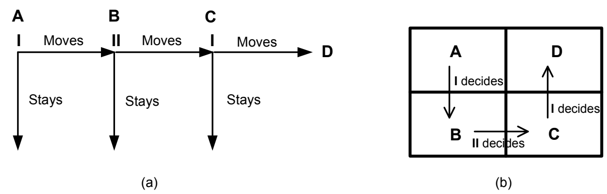

Figure 1: (a) **A game tree representation (extensive form) of our two-player game. Because of its particular structure, such games are also sometimes called Centipede games. States of the game are indicated by the letters, A, B, C and D.** (b) **Arrows denote the progression of play in the game. An action of *move* by each player causes a transition of the state of the game.**

In order to decide whether to move or stay at state **A**, a rational player *I* must reason about whether player *II* will choose to move or stay at **B**. A rational player *II*'s choice in turn depends on whether player *I* will move or stay at **C**. The default action is to stay. Thus, the game lends itself naturally to recursive reasoning and the level of reasoning is governed by the height of the game tree. Player *I*'s rational choice may be computed using backward induction in a straightforward way.

Recent studies have reported that more normative reasoning is facilitated in humans by using multimodal, spatial, ecologically valid (realistic) and experience-based presentations of the underlying game structure. Such facilitation has been reported in multiple psychological phenomena such as base-rate neglect [19, 15], overconfidence [14], the confirmation bias [13], and the conjunction fallacy [14]. In all these cases, the implementation of more realistic settings with richer contexts has made reasoning significantly closer to rational. Therefore, we imposed the following *cover story* in order to instantiate the game of Fig. 1:

Player I (you) wants to gather critical information about a target. Player II (a UAV) wants to prevent you from gathering that information. In order to gather the information, you can use your best spy-grade binoculars, and you can move closer to the target. You are currently at position J in Fig. 2(a), where you are close enough to have a 60% chance of getting the information you need. If you move to position K, you will still have a 60% chance of getting the information. If you move to position L, you increase your chances to 100%. You cannot move directly from position J to position L, and you cannot move backwards (L to K, K to J, or L to J). Player II has equipment to jam your signal, which completely destroys any information you have obtained, if it is deployed successfully. Player II is currently at position X, and it can move to position Y, but it cannot move backwards (Y to X). If you stay at position J, you will not arouse Player II's suspicion, she will not attempt to jam your signal, and she will not move from position X.

If you move to K, then Player II must choose whether to stay at X or move to Y. If Player II stays at X, she has a 33.3% chance of jamming your signal, but if she moves to Y, she has a 66.7% chance of jamming your signal. However, if she chooses to move to Y, you can quickly move from K to L while she is moving. If you are at L and she is at Y, she has a 20% chance of jamming your signal. If you move from K to L while Player II is still at X, she has a 100% chance of jamming your signal, so you should not do that.

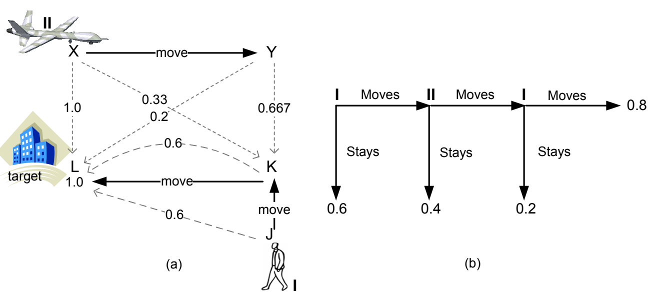

Figure 2: (a) **A spatial visualization of the game where player** I **is a human intending to gather information about a target. Player** II **is a human-controlled UAV aiming to hinder** I **from gathering the critical information. The dashed arrows and probabilities indicate the chances of** I **gathering information or** II **hindering its access.** (b) **Centipede representation of our game with the outcomes as the probabilities of success of player** I**. It is a fixed-sum game and the remaining probability is the chance of success of player** II **(failure of player** I**).**

This scenario is accompanied with Fig. 2 for illustration. If participants have difficulty in understanding the scenario, they will be further provided with a chart shown in Table 1 in which the posterior probabilities of success are clearly given. In order to succeed, player I must both obtain the information and not have the signal jammed. So the overall probability of success is: Pr(obtaining info)×[1-Pr(jamming signal)].

Notice from Fig. 2 that a rational player I will choose to stay. This is because if I chooses to move, player II will choose to stay with an overall chance of 0.6 of hindering access. A move by player II is not rational because player I will then choose to move as well with the probability of success for II being only 0.2.

Player I's position	Player II's position	I's chance of obtaining info	II's chance of jamming signal	I's overall chance of success
J	X	60%	0% (doesn't try)	60%
J	Y	wouldn't happen – II doesn't move if I doesn't go from J to K		
K	X	60%	33.3%	40%
K	Y	60%	66.7%	20%
L	X	100%	100%	0% (so don't try it!)
L	Y	100%	20%	80%

Table 1: Chart showing the various probabilities for players I **and** II**.**

3.2 Participants

A total of 162 subjects participated in the study. The participants were undergraduate students enrolled in lower-level Psychology courses at the University of Georgia. In addition to receiving performance-contingent monetary incentives, which we describe below, the participants received partial course credit.

All participants gave informed consent for their participation prior to admission into the study. They were appropriately debriefed at the conclusion of the study.

3.3 Methodology

3.3.1 Opponent Models

In order to test different levels of recursive reasoning, we designed the computer opponent (player II) to play a game in two ways: (i) If player I chooses to move, II decides on its action by simply choosing between the outcomes at states **B** and stay, and **C** with a default of stay in Fig. 1(b) rationally. Therefore, II is a zero-level player and we call it *myopic* (see Fig. 3(a)). (ii) If player I chooses to move, the opponent decides on its action by reasoning what player I will do rationally. Based on the action of I, player II will select an action that maximizes its outcomes. Thus, player II is a first-level player, and we call it *predictive* (see Fig. 3(b)).

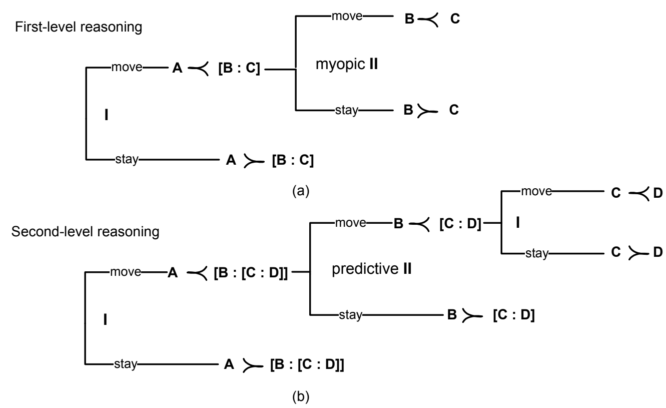

Figure 3: (a) **A myopic player** II **decides on its action by comparing the payoff at state B with that at C. Here, B** $\prec$ **C denotes a preference of C over B for the player whose turn it is to play and B : C denotes the rational choice by the appropriate player between actions leading to states B and C. Thus, player** I **exhibits first-level reasoning.** (b) **If player** II **is predictive, it reasons about** I**'s actions. Player** I **then exhibits second-level reasoning in deciding its action at state A.**

To illustrate, in the game of Fig. 2 if player I decides to move, a myopic player II will move to obtain a probability of success of 0.8, while a predictive II will choose to stay because it thinks that player I will choose to move from **C** to **D**, if it moved. By choosing to stay, II will obtain an outcome of 0.6 in comparison to 0.2 if it moves.

3.3.2 Payoff Structure

Notice that the rational choice of players in the game of Fig. 1 depends on the preferential ordering of states of the game rather than actual values. Let $a \prec b$ indicate that the player whose turn it is to play prefers state b over a, and because the game is purely competitive, the other player prefers state a over b. Games that exhibit a preference ordering of **D** $\prec$ **C** $\prec$ **B** $\prec$ **A** and **A** $\prec$ **B** $\prec$ **C** $\prec$ **D** for player I are trivial because player I will always opt to stay in the former case and move in the latter case, regardless of how II plays. Furthermore, consider the ordering **C** $\prec$ **A** $\prec$ **B** $\prec$ **D** for player I. A myopic opponent will choose to move while a predictive opponent will stay. However, in both these cases player I will choose to move. Thus, games whose states display a preferential ordering of the type mentioned previously are not diagnostic – regardless of whether player I thinks that opponent is myopic or is predictive, I will select the same action precluding a diagnosis of

I's level of recursive reasoning. Of all the 24 distinct preferential orderings among states that are possible, only one is diagnostic: **C** $\prec$ **B** $\prec$ **A** $\prec$ **D**. For this ordering, player *I* will move if it thinks that the opponent is myopic, otherwise *I* will stay if the opponent is thought to be predictive. We point out that the game in Fig. 2 follows this preference ordering.

3.3.3 *Design of Task*

Batches of participants played the game on computer terminals with each batch having an even number of players. Each batch was divided into two groups and members of the two groups were sent to different rooms. This was done to create the illusion that each subject was playing against another, although the opponent was in reality a computer program. This deception was revealed to the subjects during debriefing.

Each subject experienced an initial *training phase* of at least 15 games that were trivial or those in which a myopic or predictive opponent behaved identically. These games served to acquaint the participants with the rules and goal of the task without unduly biasing them about the behavior of the opponent. Therefore, these games have no effect on the initial model of the opponent that participants may have. Participants who failed to choose the rational actions in any of the previous 5 games after the 15-game training phase continued with new training games until they met the criterion of no rationality errors in the 5 most recent games. Those who failed to meet this criterion after 40 total training games did not advance to the test phase, and were removed from the study.

In the *test phase*, each subject experienced 40 games instantiated with outcome probabilities that exhibited the diagnostic preferential ordering of **C** $\prec$ **B** $\prec$ **A** $\prec$ **D** for player *I*. The 40 critical games were divided into 4 blocks of 10 games each. In order to avoid subjects developing a mental set, we interspersed these games with 40 that exhibited the orderings, **C** $\prec$ **A** $\prec$ **B** $\prec$ **D** and **D** $\prec$ **B** $\prec$ **A** $\prec$ **C**. The latter games not only serve to distract the participants but also function as "catch" trials allowing us to identify participants who may not be attending to the games.

Approximately half the participants played against myopic opponents while the remaining played against predictive ones. In each category, approximately half of the participants were presented with just the Centipede representation of the games with probabilistic payoffs and no cover story, which we label as the *abstract* version. Remaining participants in the category were presented with the UAV cover story and the associated picture in Fig. 2, including the Centipede representation. We label this as the *realistic* version. About half of all participants also experienced a screen asking them what they thought the opponent would play and their confidence in the prediction, for some of the games.

Participants received a monetary incentive of 50 cents for every correct action that they chose in a game. This resulted in an average payout of approximately $30 per participant.

3.4 Results and Discussion

Our study spanned a period of *three months* from September through November 2008. We report the results of this study below.

3.4.1 *Training Phase*

As mentioned before, each of the 162 human subjects initially played a series of 15 games in order to get acquainted with the fixed-sum and complete information structure, and objectives of the task at hand. After this initial phase, participants who continued to exhibit errors in any of the games up to 40 total games were eliminated. 26 participants did not progress further in the study. These participants either failed to understand how the game is played or exhibited excessive irrational behavior, which would have affected the validity of the results of this study.

3.4.2 *Test Phase*

Of the 136 participants (70 female) who completed the test phase, we show the numbers that experienced myopic or predictive opponents and abstract or realistic versions of the games, in Table 2.

Structure	myopic abstract	myopic realistic	predictive abstract	predictive realistic
No. of subjects	**37**	**30**	**37**	**32**

Table 2: The numbers of participants that experienced each of the 4 different types of tasks. The numbers differ from each other because of eliminations in the training phase.

Participants in each of the 4 groups were presented with 40 instances of the particular game type whose payoff structure is diagnostic. For the sake of analysis, we assembled 4 *test blocks* each comprising 10 games. For each participant, we measured the fraction of times that the subject played accurately in each test block. We define an *accurate choice* as the action choice which is rational given the type of opponent. For example, in the game of Fig. 2, the accurate choice for player *I*, if the opponent is myopic, is to move. On the other hand, if the opponent is predictive, the accurate choice for *I* is to stay.

Because opponents types are fixed and participants experience 40 games, they have the opportunity to learn how their opponent might be playing the games. Consequently, participants may gradually make more accurate choices over time. Participants were deemed to have learnt the opponent's model at the game beyond which performance was always statistically significantly better than chance, as measured by a binomial test at the 0.05 level and one-tailed. This implies making no more than one inaccurate choice in any block of 10 games. (For this purpose, blocks were defined by a moving window of 10 games, not the fixed blocks used in other analyses.)

In Fig. $4(a)$, we show the mean proportion of accurate choices across all participants in each of the 4 groups. Two group-level findings are evident from the results in Fig. $4(a)$: First, the mean proportion of accurate choices is significantly higher when the opponent is predictive as compared to when it is myopic. This is further evident from Fig. $4(b)$ where we show the mean proportions *marginalized* over the abstract and realistic versions of the tasks. Student t-tests with p-values < 0.0001 confirm that participants playing against predictive opponents have statistically significant higher proportions of accuracy compared to myopic opponents across all test blocks.

The higher proportions of accurate choice when the opponent is predictive in conjunction with the lower proportions when the opponent is myopic implies that subjects predominantly displayed second-level reasoning when acting. They expected the opponent to reason about their subsequent play (first-level reasoning) and acted accordingly. The fact that myopic opponents did not do this resulted in their choices being inaccurate.

Second, no significant difference in the mean proportions between abstract and realistic versions of the tasks is evident across any test block from Fig. $4(a)$. This is regardless of whether the opponent is myopic or predictive. This observation is further evident in Fig. $4(c)$ which shows the mean proportions for abstract and realistic versions marginalized over myopic and predictive opponents. Student t-tests with very low p-values revealed no statistical signif-

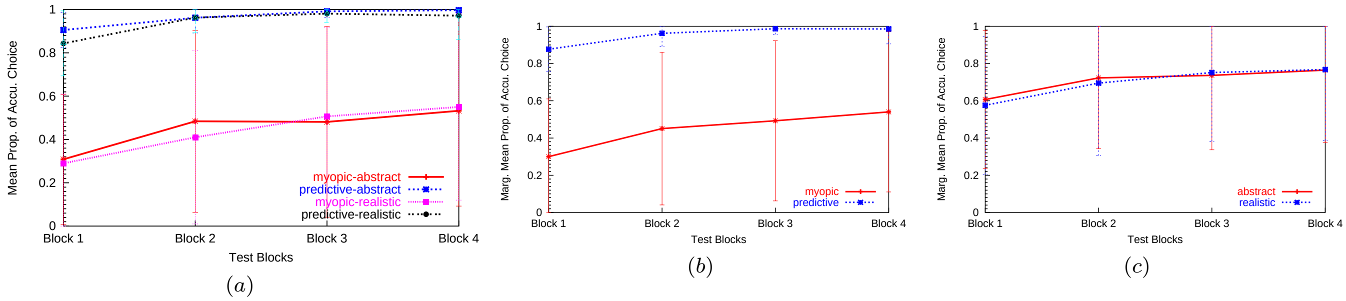

Figure 4: (a) Mean proportion of accurate choices of the participants for all conditions across test blocks. Notice that subjects generally expected their opponents to play at first level far more than at zero level. Mean proportions marginalized over (b) abstract and realistic versions of the task, and (c) over myopic and predictive opponents. The difference in the latter is not statistically significant.

icance in the difference between the proportions, either overall or in any test block.

The lack of any significant difference in the accuracy of the choices seems to suggest that our cover story neither confounded the participants nor clarified it further in an intuitive sense. We speculate that the indifference is due to, (i) the Centipede representation though abstract being sufficiently clear to facilitate understanding of this simple game, and thus (ii) subjects playing the games with high accuracy leaving little room for improvement, at least in the predictive groups.

Notice from Fig. 4(a) that the mean proportion of accurate choice improves over successive test blocks in all groups. Many participants in the predictive groups learnt in fewer than 10 games, by making no inaccurate choices in the first 7 games, and no more than one in the first 10 games. Forty-one participants, of which 40 were in the myopic conditions, never achieved learning by our standard; they were then assigned a value of 40 for the number of games to learning. The average number of games to learning was 10 for the predictive group and 31.7 for the myopic group (for the difference, Student t-tests reveal $p < 0.0001$).

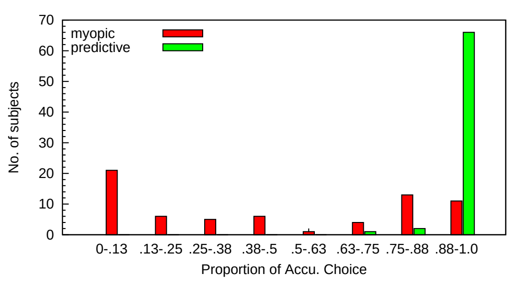

Figure 5: Count of participants who played myopic or predictive opponents and grouped according to different proportions of accurate choice.

Finally, in Fig. 5, we detail the number of participants whose actions across all games fell into different bins of proportion of accurate choice. Fig. 5 reveals that about 83% of the 67 participants who played against a myopic opponent had proportions less than 0.875. In comparison, only about 4% of the 69 participants who played a predictive opponent exhibited such proportions of accurate choice. Consequently, these results conclusively reveal that our subjects predominantly played reasonably and as though they expected the opponents to be predictive, and thus generally reasoned at a level higher than had been previously observed.

4. DISCUSSION

Multiple previous investigations into spontaneous recursive reasoning by individuals while playing strategic games have shown that humans *generally* do not think that opponents will think recursively while playing. Therefore, they themselves reason at only a single level of nesting. This has been attributed to the bounded rationality and limited mental processes of humans. Using a strategic game employed in one such previous study but made more competitive by incorporating fixed-sum payoffs and tangible incentives, we have shown that humans exhibit a higher level of recursive thinking when playing this game. Thus, we have demonstrated that in some settings humans *generally* reason at higher levels of recursion, and therefore exhibit more rational behavior. Consequently, our experiment opens up avenues for identifying settings where humans typically exhibit other forms of strategic sophistication such as simultaneous deeper and longer term thinking.

Although there is psychological evidence to suggest that intuitive cover stories induce human behavior closer to rational action, our cover story of UAV reconnaissance neither reduced nor improved the accuracy of the results. However, we do not claim that the previously perceived effect is not real, only that our particular cover story did not evoke it. We suspect that it may make a positive difference if the game is more sophisticated such as an extended Centipede game requiring three or four levels of recursive thinking.

An alternate hypothesis for explaining our results could be that many of the subjects solved the games completely using backward induction (ie. minimax). However, we do not believe this to be the case because, (a) participants are implicitly encouraged to think about opponent models, (b) Hedden and Zhang's result of predominant first-level reasoning in an identically-structured game precludes the use of backward induction. Others have noted that backward induction fails to explain a variety of human reasoning and decision-making behavior [23], and that it is not robust particularly in gaming situations where parity and certainty do not exist [3]. Our game does not display parity because player I has greater control over the outcome of the game, and (c) the subjects were not explicitly trained in sophisticated game-theoretic techniques such as backward induction or minimax for solving games. Indeed, our exit questionnaire revealed that participants predominantly reasoned about the opponent's thinking.

Acknowledgment

This research is supported in part by grant #FA9550-08-1-0429 from AFOSR.

5. REFERENCES

[1] R. J. Aumann. Interactive epistemology i: Knowledge. *International Journal of Game Theory*, 28:263–300, 1999.

[2] R. J. Aumann. Interactive epistemology ii: Probability. *International Journal of Game Theory*, 28:301–314, 1999.

[3] S. J. Brams and D. M. Kilgour. Backward induction is not robust: the parity problem and the uncertainty problem. *Theory and Decision*, 45:263–289, 1998.

[4] A. Brandenburger. The power of paradox: Some recent developments in interactive epistemology. *International Journal of Game Theory*, 35:465–492, 2007.

[5] A. Brandenburger and E. Dekel. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59:189–198, 1993.

[6] C. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, 2003.

[7] A. Colman. *Game Theory and Experimental Games*. Oxford: Pergamon Press, 1982.

[8] P. Doshi. On the role of interactive epistemology in multiagent planning. In *International Conference on AI and Pattern Recognition*, pages 208 – 213, 2007.

[9] R. Dunbar. Theory of mind and evolution of language. In *Approaches to the Evolution of Language*. Cambridge University Press, 1998.

[10] R. Fagin, J. Geanakoplos, J. Halpern, and M. Vardi. The hierarchical approach to modeling knowledge and common knowledge. *International Journal of Game Theory*, 28, 1999.

[11] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.

[12] S. Ficici and A. Pfeffer. Modeling how humans reason about others with partial information. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 315–322, 2008.

[13] L. Fiddick, L. Cosmides, and J. Tooby. No interpretation without representation: The role of domain-specific representations and inferences in the wason selection task. *Cognition*, 77:1–79, 2000.

[14] G. Gigerenzer. How to make cognitive illusions disappear: Beyond "heuristics and biases". *European Review of Social Psychology*, 2:83–115, 1991.

[15] G. Gigerenzer and U. Hoffrage. How to improve bayesian reasoning without instructions: frequency formats. *Psychological Review*, 102:684–704, 1995.

[16] P. Gmytrasiewicz and P. Doshi. A framework for sequential planning in multiagent settings. *Journal of Artificial Intelligence Research*, 24:49–79, 2005.

[17] P. J. Gmytrasiewicz and E. H. Durfee. A rigorous, operational formalization of recursive modeling. In *International Conference on Multi-Agent Systems*, pages 125–132, 1995.

[18] P. J. Gmytrasiewicz and E. H. Durfee. Rational interaction in multiagent environments: Coordination. *Journal of Autonomous Agents and Multi-Agent Systems*, 3:319–350, 2000.

[19] A. Goodie and E. Fantino. Learning to commit or avoid the base-rate error. *Nature*, 380:247–249, 1996.

[20] J. C. Harsanyi. Games with incomplete information played by bayesian players. *Management Science*, 14(3):159–182, 1967.

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

[22] A. Heifetz and D. Samet. Topology-free typology of beliefs. *Journal of Economic Theory*, 82:324–341, 1998.

[23] R. D. McKelvey and T. R. Palfrey. An experimental study of the centipede game. *Econometrica*, 60:803–836, 1992.

[24] J. Mertens and S. Zamir. Formulation of bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14:1–29, 1985.

[25] D. Stahl and P. Wilson. On player's models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10:218–254, 1995.