Human strategic reasoning in dynamic games: Experiments, logics, cognitive models

Sujata Ghosh[1], Tamoghna Halder[2], Khyati Sharma[3] and Rineke Verbrugge[4]

[1] Indian Statistical Institute, Chennai, sujata@isichennai.res.in
[2] Indian Statistical Institute, Kolkata, thaldera@gmail.com
[3] Dunia Finance LLC, Dubai, khyati.sharma27@gmail.com
[4] University of Groningen, The Netherlands, L.C.Verbrugge@rug.nl

Abstract. This article provides a three-way interaction between experiments, logic and cognitive modelling so as to bring out a shared perspective among these diverse areas, aiming towards better understanding and better modelling of human strategic reasoning in dynamic games.

1 Introduction

How suitable are idealized formal models of social reasoning processes with respect to the nuances of the real world? In particular, do these formal methods represent human strategic reasoning satisfactorily or should we instead concentrate on empirical studies and models based on those empirical data? Ghosh, Meijering and Verbrugge [6] made an effort to bridge the gap between logical and cognitive treatments of strategic reasoning in dynamic games. They proposed to combine empirical studies, formal modeling and cognitive modeling to study human strategic reasoning. In their words, "rather than thinking about logic and cognitive modeling as completely separate ways of modeling, we consider them to be complementary and investigate how they can aid one another to bring about a more meaningful model of real-life scenarios". In the current article, we apply this combination of methods to the question to what extent people use backward induction or forward induction in dynamic games.

Backward and forward induction reasoning Backward Induction (BI) is the textbook approach for solving extensive-form games with perfect information. In generic games without payoff ties, BI yields the unique subgame perfect equilibrium. The assumptions underpinning BI are that all players commonly believe in everybody's future rationality, no matter how irrational players' past behaviour has already proven. See [15,18] for more details.

In Forward Induction (FI) reasoning, on the other hand, a player tries to rationalize the opponent's past behaviour in order to assess his future moves. Thus, in a subgame where no strategy of the opponent is consistent with common knowledge of rationality *and* his past behaviour, the player may still rationalize the opponent's past behaviour by attributing to him a strategy which is optimal against a presumed *suboptimal* strategy of hers, or by attributing to him a strategy which is optimal vis-a-vis a *rational* strategy of hers, which is only optimal against a suboptimal strategy of *his*. If the player pursues this rationalizing reasoning to the highest extent possible [2] and reacts accordingly, she ends up choosing what is called an *Extensive-Form Rationalizable* (EFR) strategy [17] (see also [18,16,5]). Thus EFR strategies are based on FI reasoning, and in the following we use the terms EFR and FI synonymously.

There have been extensive debates among game theorists and logicians about the merits of backward induction. Experimental economists and psychologists have shown that human subjects do not always follow the backward induction strategy in large centipede games [10,14]. Recently, based on an eye-tracking study and complexity considerations, it turned out that even when human subjects produce the outwardly correct 'backward induction answer' in smaller games, they may use a different internal reasoning strategy to achieve it [13,3]. To investigate human reasoning strategies, Ghosh, Meijering and Verbrugge [6] presented a formal language to represent strategies on a finer-grained level than was possible before. The language and its semantics helped to precisely distinguish different cognitive reasoning strategies, that can then be tested on the basis of computational cognitive models and experiments with human subjects. The syntactic framework of the formal system provided a generic way of constructing computational cognitive models of the participants of a 'marble drop' game.

Aims of this article Ghosh, Heifetz and Verbrugge [5] conducted a game-theoretic experiment that involves a participant's expectations about the opponent's reasoning strategies, that may in turn depend on expectations about the participant's reasoning. It deals with the following question: In a dynamic game of perfect information, are people inclined to do forward induction reasoning (i.e. show EFR behaviour)? In the current work, we extend our aim of bridging formal and empirical studies to this question from behavioural game theory, utilizing the experimental findings from [5]. The main new elements of this work with respect to [6,5,8] are as follows:

- We study robustness of the findings of [5], to alleviate concern that different participants might follow a variety of reasoning patterns. Thus, more grounding is given to the outcomes, which is used for formal modelling.
- Unlike the eye-tracking studies used in [13,6], the experiment which forms the backbone of this paper includes participants' verbal comments regarding the reasoning they applied to perform their actions (see [8]), which made it possible to introduce agents' beliefs about their opponents' moves and beliefs in the logical language. We conjecture that this language is more succinct than the language proposed in [6] in describing strategic reasoning, which in turn may lead to a more efficient modelling.

In what follows, we briefly recall Ghosh and colleagues' recent experiment on forward induction [5,8], report a robustness study of the findings of the experiment, and extend the language introduced in [6] to describe players' reasoning strategies, adding a belief operator to reflect players' expectations. Finally, we sketch how strategy-formulas in this extended language can be turned into computational cognitive models that help to distinguish what is going on in people's minds when they play dynamic games of perfect information.

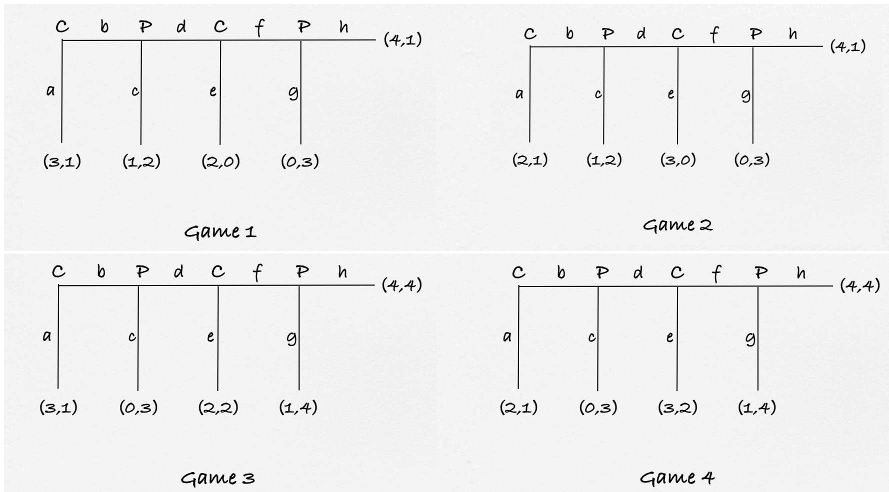

Fig. 1. Collection of the main games used in the experiment. The ordered pairs at the leaves represent pay-offs for the computer (C) and the participant (P), respectively.

2 An experimental study: do people use FI?

We provide a brief summary of the experimental games and the experimental procedure underlying the current work. The experiment (previously reported in [5]) was designed to tackle the question whether people are inclined to use forward induction (FI) reasoning when they play dynamic perfect information games. The main interest was to examine participants' behaviour following a deviation from BI behaviour by their opponent right at the beginning of the game; for details, see [5,8].

The games that were used in the experiment are given in Figures 1 and 2. In these two-player games, the players play alternately. Let C denote the computer and P the participant. In the first four games (Figure 1), the computer plays first, followed by the participant. The players control two decision nodes each. In the last two games (Figure 2), which are truncated versions of two of the games of Figure 1, the participant moves first.

To explicate the difference between BI and EFR behaviour consider game 1, one of the experimental games (cf. Figure 1). Here, the unique Backward Induction (BI) strategies for player C and player P are $a; e$ and $c; g$, respectively, which indicate that the game will end at the first node, going down. In contrast, EFR would proceed as follows, starting from the scenario in which the game reaches the first decision node of P. Among the two strategies of player C that are compatible with this event, namely $b; e$ and $b; f$, only the latter is rational for player C. This is because of the fact that $b; e$ is dominated by $a; e$, while $b; f$ is optimal for player C if she believes that player P will play $d; h$ with a high enough probability. Attributing to player C the strategy $b; f$ is thus player P's best way to rationalize player C's choice of b, and in reply, $d; g$ is player P's best response to $b; f$. Thus, the unique Extensive-Form Rationalizable (EFR, [17]) strategy

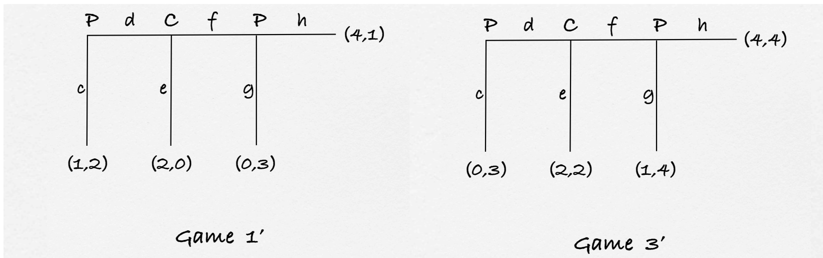

Fig. 2. Truncated versions of Game 1 and Game 3. The ordered pairs at the leaves represent pay-offs for C and P, respectively.

(an *FI* strategy) of player P is $d; g$, which is distinct from his BI strategy $c; g$. For a detailed discussion on BI and EFR strategies in games $2, 3, 4, 1', 3'$, see [5].

The experiment was conducted at the Institute of Artificial Intelligence at the University of Groningen, the Netherlands. A group of 50 Bachelor's and Master's students from different disciplines took part. They had little or no knowledge of game theory, so as to ensure that neither backward induction nor forward induction was already known to them.[5] The participants played the finite perfect-information games in a graphical interface on the computer screen (cf. Figure 3). In each case, the opponent was the computer which had been programmed to play according to plans that were best responses to some plan of the participant, and this was told to the participants.

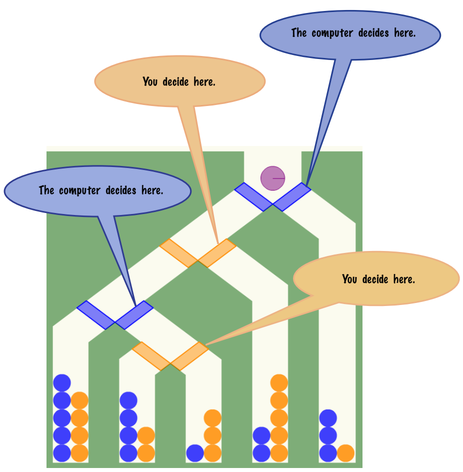

Fig. 3. Graphical interface for the participants. The computer controls the blue trapdoors and acquires blue marbles (represented as dark grey in a black and white print) as pay-offs, while the participant controls the orange trapdoors and acquires orange marbles (light grey in a black and white print) as pay-offs.

After 14 practice games, each participant played 48 experimental games. There were 8 rounds, each comprised of 6 games as described above. Differ-

[5] The candidate participants were asked about their educational details. Two students who had followed a course on game theory were excluded.

ent graphical representations of the same game were used in different rounds. Participants earned 10-15 euros for participation, depending on points earned.

At the end of the experiment, each participant was asked the following question: 'When you made your choices in these games, what did you think about the ways the computer would move when it was about to play next?' The participant needed to describe in his own words, the plan he thought was followed by the computer on its next move after the participant's initial choice. We used these answers to classify various strategic reasoning processes applied by the participants while playing the experimental games.

To analyse whether participants P played FI strategies in the games described in figures 1 and 2, we can formulate the following *hypothesis* (see [5] for an explanation) concerning the participant's choice in his first decision node (if reached in games 1, 2, 3, 4, and in all rounds of games $1'$ and $3'$):

> "d will be played most often in game 3, less so in game 1, even less in games $3'$ and 4, least often in games $1'$ and 2", which we henceforth abbreviate as "$d : 3 > 1 > 3', 4 > 1', 2$."

In games 1 and 3, d is the only EFR move; in games $1'$ and 2, d is neither a BI nor an EFR move; and in games $3'$ and 4, both c and d are EFR moves. Moreover, in game 3, reaching the first decision node is compatible with common knowledge of rationality.

Ghosh et al. [5] found that in the aggregate, participants were indeed more likely to make decisions in accordance with their best-rationalization EFR conjecture, i.e., consistent with FI reasoning. For a detailed study and a discussion of some alternative explanations of the results, see [5,8]. Our main concern in the current paper is how we can construct cognitive models based on the experimental findings and how logic can play a role in such construction. To justify our aim, we first investigate the robustness of the results of [5] based on the available group-divisions.

2.1 Robustness: different results for different groups?

We segregated the participants in terms of gender and discipline and went on to test the hypothesis over the different groups formed by segregation.[6]

Segregation by gender The available data on the behaviour of participants at their first decision node in the six games were divided into two groups: male and female. Overall, 40 men and 10 women had participated in the experiment reported in [5,8]. We studied the choices made by participants belonging to the two groups.[7] For the hypothesis, we have the following, very similar to the results reported in [5]:

- **Male** $d : 3, 3' > 4 > 1' > 1 > 2$
- **Female** $d : 3, 3' > 4 > 1', 1 > 2$

[6] Because of little variance among participants, we did not segregate by age.

[7] The results are based on one sample and two sample proportion tests.

As to individual games, the tests revealed the following behaviour. We use the notations $i \sim j$ to denote that options i and j are chosen equally often, and $i > j$ to denote that i is chosen more often than j. The null hypothesis was that c and d were chosen equally often at the first decision node:

– Game 1: $c > d$ (male, female).
– Game 2: $c > d$ (male, female).
– Game 3: $d > c$ (male, female).
– Game 4: $d \sim c$ (male), $d > c$ (female).
– Game 1′: $c > d$ (male, female).
– Game 3′: $d > c$ (male, female).

Segregation by discipline For this study, the data on 50 participants was separated into three broad groups based on the nature of the study fields of the participants:

Artificial Intelligence (AI): artificial intelligence and human-machine communication (27 students);
Behavioural and Social Sciences (BSS): accountancy, economics and business economics, human resource management, international relations, law and business economics, and psychology (10 students);
Exact Sciences (ES): biology, biomedical sciences, drug innovation, computer science, mathematics, and physics (13 students).

Similar statistical analysis was done over the choices made by the participants belonging to the three groups. We summarize the results for the hypothesis:

– **AI** $d : 3, 3' > 4 > 1', 1 > 2$
– **BSS** $d : 3, 3' > 4 > 1', 1 > 2$
– **ES** $d : 3, 3' > 4 > 1', 1 > 2$

For the hypotheses on the individual games:

– Game 1: $c > d$ (AI, BSS), $d \sim c$ (ES).
– Game 2: $c > d$ (AI, BSS, ES).
– Game 3: $d > c$ (AI, BSS, ES).
– Game 4: $d \sim c$ (BSS, ES), $d > c$ (AI).
– Game 1′: $c > d$ (AI, BSS), $d \sim c$ (ES).
– Game 3′: $d > c$ (AI, BSS, ES).

The statistical analyses based on gender and discipline suggest that the results mentioned in Section 2 about participants' behaviour at their first decision node are robust. We only found minor variations corresponding to certain groups.

3 A language for strategies

In the line of [6], we propose a logical language specifying strategies of players. Our motivation for introducing this logical framework is to build a pathway from empirical to cognitive modelling studies. A detailed formal study of this framework regarding its expressive power and axiomatics is left for future work.

This framework uses empirical studies to provide insights into cognitive models of human strategic reasoning as performed during the experiment discussed in Section 2. The main idea is to use the logical syntax to express the different reasoning procedures as performed and conveyed by the participants and use these formulas to systematically build up reasoning rules of computational cognitive models of strategic reasoning.

A novel part of the proposed language is that we add an explicit notion of belief to the language proposed in [6] in order to describe participants' expectations regarding future moves of the computer. This belief operator is parametrized by both players and nodes of the game tree so that the possible expectations of players at each of their nodes can be expressed within the language itself. The whole point is to explicate the human reasoning process, therefore the participants' beliefs and expectations need to come to the fore. Such expectations formed an essential part of the current experimental study. We first build a syntax for game trees (cf. [19,7]). Let N denote a finite set of players and let Σ denote a countable set of actions.

Syntax for extensive form game trees Let *Nodes* be a countable set. The syntax for specifying finite extensive form game trees is given by:

$$\mathbb{G}(Nodes) := (i,x) \mid \Sigma_{a_m \in J}((i,x), a_m, t_{a_m})$$

where $i \in N$, $x \in Nodes$, $J(\text{finite}) \subseteq \Sigma$, and $t_{a_m} \in \mathbb{G}(Nodes)$.

Given $h \in \mathbb{G}(Nodes)$, we define the tree T_h generated by h inductively as follows (see Figure 4 for an example):

- $h = (i,x)$: $T_h = (S_h, \Rightarrow_h, \widehat{\lambda}_h, s_x)$ where $S_h = \{s_x\}$, $\widehat{\lambda}_h(s_x) = i$.
- $h = ((i,x), a_1, t_{a_1}) + \cdots + ((i,x), a_k, t_{a_k})$: Inductively we have trees $T_1, \ldots T_k$ where for $j : 1 \leqslant j \leqslant k$, $T_j = (S_j, \Rightarrow_j, \widehat{\lambda}_j, s_{j,0})$.
  Define $T_h = (S_h, \Rightarrow_h, \widehat{\lambda}_h, s_x)$ where
  - $S_h = \{s_x\} \cup S_{T_1} \cup \ldots \cup S_{T_k}$;
  - $\widehat{\lambda}_h(s_x) = i$ and for all j, for all $s \in S_{T_j}$, $\widehat{\lambda}_h(s) = \widehat{\lambda}_j(s)$;
  - $\Rightarrow_h = \bigcup_{j:1 \leqslant j \leqslant k}(\{(s_x, a_j, s_{j,0})\} \cup \Rightarrow_j)$.

Given $h \in \mathbb{G}(Nodes)$, let $Nodes(h)$ denote the set of distinct pairs (i,x) that occur in the expression of h.

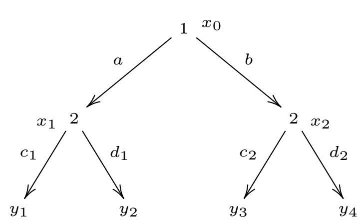

Fig. 4. Extensive form game tree. The nodes are labelled with turns of players and the edges with the actions. The syntactic representation of this tree can be given by:
$h = ((1, x_0), a, t_1) + ((1, x_0), b, t_2)$, where
$t_1 = ((2, x_1), c_1, (2, y_1)) + ((2, x_1), d_1, (2, y_2))$;
$t_2 = ((2, x_2), c_2, (2, y_3)) + ((2, x_2), d_2, (2, y_4))$.

3.1 Strategy specifications

A syntax for specifying partial strategies and their compositions in a structural manner involving simultaneous recursion has been used in [6] to describe empirical reasoning of participants involved in a game experiment in a dynamic game called 'marble drop' [12,11], as demonstrated by an eye-tracking study [13]. The main case specifies, for a player, which conditions she tests before making a move. In what follows, the pre-condition for a move depends on observables that hold at the current game position, some belief conditions, as well as some simple finite past-time conditions and some finite look-ahead that each player can perform in terms of the structure of the game tree. Both the past-time and future conditions may involve some strategies that were or could be enforced by the players. These pre-conditions are given by the syntax defined below.

For any countable set X, let $BPF(X)$ (the boolean, past and future combinations of the members of X) be sets of formulas given by the following syntax:

$$BPF(X) := x \in X \mid \neg\psi \mid \psi_1 \vee \psi_2 \mid \langle a^+ \rangle\psi \mid \langle a^- \rangle\psi,$$

where $a \in \Sigma$, a countable set of actions.

Formulas in $BPF(X)$ can be read as usual in a dynamic logic framework and are interpreted at game positions. The formula $\langle a^+ \rangle\psi$ (respectively, $\langle a^- \rangle\psi$) refers to one step in the future (respectively, past). It asserts the existence of an a edge after (respectively, before) which ψ holds. Note that future (past) time assertions up to any bounded depth can be coded by iteration of the corresponding constructs. The 'time free' fragment of $BPF(X)$ is formed by the boolean formulas over X. We denote this fragment by $Bool(X)$.

For each $h \in \mathbb{G}(Nodes)$ and $(i, x) \in Nodes(h)$, we now add a new operator $\mathbb{B}_h^{(i,x)}$ to the syntax of $BPF(X)$ to form the set of formulas $BPF_b(X)$. The formula $\mathbb{B}_h^{(i,x)}\psi$ can be read as "in the game tree h, player i *believes* at node x that ψ holds". One might feel that it is not elegant that the belief operator is parametrized by the nodes of the tree, however, our main aim is not to propose a logic for the sake of its nice properties, but to have a logical language that can be used suitably for constructing computational cognitive models corresponding to participants' strategic reasoning.

Syntax Let $P^i = \{p_0^i, p_1^i, \ldots\}$ be a countable set of observables for $i \in N$ and $P = \bigcup_{i \in N} P^i$. To this set of observables we add two kinds of propositional variables $(u_i = q_i)$ to denote 'player i's utility (or payoff) is q_i' and $(r \leqslant q)$ to denote that 'the rational number r is less than or equal to the rational number q'[8] The syntax of strategy specifications is given by:

$$Strat^i(P^i) := [\psi \mapsto a]^i \mid \eta_1 + \eta_2 \mid \eta_1 \cdot \eta_2,$$

where $\psi \in BPF_b(P^i)$. For a detailed explanation see [6]. The basic idea is to use the above constructs to specify properties of strategies as well as to combine them to describe a play of the game. For instance, the interpretation of a player i's specification $[p \mapsto a]^i$ where $p \in P^i$, is to choose move a at every game position belonging to player i where p holds. At positions where p does not hold, the strategy is allowed to choose any enabled move. The strategy specification $\eta_1 + \eta_2$ says that the strategy of player i conforms to the specification η_1 or η_2. The construct $\eta_1 \cdot \eta_2$ says that the strategy conforms to specifications η_1 and η_2.

[8] as in [6] and inspired by [4].

Semantics We consider perfect information games with belief structures as models. The idea is very similar to that of temporal belief revision frame presented in [4]. Let $M = (T, \{\longrightarrow_i^x\}, V)$ with $T = (S, \Rightarrow, s_0, \widehat{\lambda}, \mathcal{U})$, where $(S, \Rightarrow, s_0, \widehat{\lambda})$ is an extensive form game tree, $\mathcal{U} : \mathit{frontier}(T) \times N \to \mathbb{Q}$ is a utility function. Here, $\mathit{frontier}(T)$ denotes the leaf nodes of the tree T. For each $s_x \in S$ with $\widehat{\lambda}(s_x) = i$, we have a binary relation $\longrightarrow_i^x$ over S (cf. the connection between h and T_h presented above). Finally, $V : S \to 2^P$ is a valuation function. The truth value of a formula $\psi \in BPF_b(P)$ at the state s, denoted $M, s \models \psi$, is defined as follows:

- $M, s \models p$ iff $p \in V(s)$.
- $M, s \models \neg\psi$ iff $M, s \not\models \psi$.
- $M, s \models \psi_1 \vee \psi_2$ iff $M, s \models \psi_1$ or $M, s \models \psi_2$.
- $M, s \models \langle a^+ \rangle\psi$ iff there exists an s' such that $s \overset{a}{\Rightarrow} s'$ and $M, s' \models \psi$.
- $M, s \models \langle a^- \rangle\psi$ iff there exists an s' such that $s' \overset{a}{\Rightarrow} s$ and $M, s' \models \psi$.
- $M, s \models \mathbb{B}_h^{(i,x)}\psi$ iff the underlying game tree of T_M is the same as T_h and for all s' such that $s \longrightarrow_i^x s'$, $s' \models \psi$.

The truth definitions for the new propositions are as follows:

- $M, s \models (u_i = q_i)$ iff $\mathcal{U}(s, i) = q_i$.
- $M, s \models (r \leqslant q)$ iff $r \leqslant q$, where r, q are rational numbers.

Strategy specifications are interpreted on strategy trees of T. We also assume the presence of two special propositions $\mathbf{turn}_1$ and $\mathbf{turn}_2$ that specify which player's turn it is to move, i.e. the valuation function satisfies the property

- for all $i \in N$, $\mathbf{turn}_i \in V(s)$ iff $\widehat{\lambda}(s) = i$.

One more special proposition **root** is assumed to indicate the root of the game tree, that is the starting node of the game. The valuation function satisfies the property

- $\mathbf{root} \in V(s)$ iff $s = s_0$.

We recall that a strategy for player i is a function μ^i which specifies a move at every game position of the player, i.e. $\mu^i : S^i \to \Sigma$. A strategy μ can also be viewed as a subtree of T where for each node belonging to the opponent player i, there is a unique outgoing edge and for nodes belonging to player $\bar{\imath}$, every enabled move is included. A partial strategy for player i is a partial function σ^i which specifies a move at some (but not necessarily all) game positions of the player, i.e. $\sigma^i : S^i \rightharpoonup \Sigma$. A partial strategy can be viewed as a set of total strategies of the player [6].

The semantics of the strategy specifications are given as follows. Given a model M and a partial strategy specification $\eta \in \mathit{Strat}^i(P^i)$, we define a semantic function $[\![\cdot]\!]_M : \mathit{Strat}^i(P^i) \to 2^{\Omega^i(T_M)}$, where each partial strategy specification is associated with a set of total strategy trees and $\Omega^i(T)$ denotes the set of all player i strategies in the game tree T.

For any $\eta \in \mathit{Strat}^i(P^i)$, the semantic function $[\![\eta]\!]_M$ is defined inductively:

- $[\![[\psi \mapsto a]^i]\!]_M = \Upsilon \in 2^{\Omega^i(T_M)}$ satisfying: $\mu \in \Upsilon$ iff μ satisfies the condition that, if $s \in S_\mu$ is a player i node then $M, s \models \psi$ implies $\mathit{out}_\mu(s) = a$.

- $[\![\eta_1 + \eta_2]\!]_M = [\![\eta_1]\!]_M \cup [\![\eta_2]\!]_M$
- $[\![\eta_1 \cdot \eta_2]\!]_M = [\![\eta_1]\!]_M \cap [\![\eta_2]\!]_M$

Above, $out_\mu(s)$ is the unique outgoing edge in μ at s. Recall that s is a player i node and therefore by definition of a strategy for player i, there is a unique outgoing edge at s.

Before describing specific strategies found in the empirical study, we would like to focus on the new operator of belief, $\mathbb{B}_h^{(i,x)}$ proposed above. Note that this operator is considered for each node in each game. The idea is that the same player might have different beliefs at different nodes of the game. We had to introduce the syntax of the extensive form game trees to make this definition sound, otherwise we would have had to restrict our discussion to single game trees. The semantics given to the operator is entangled in both the syntax and semantics, which might create problems in finding an appropriate axiom system. A possible solution would be to introduce some generic classes of games similar to the idea of generic game boards [20], using the notion of enabled game trees [7]. This is left for future work, as well as a comparison of the expressiveness of the current language with those of existing logics of belief and strategies.

3.2 Describing specific strategies in the experimental games

Let us now express some actual reasoning processes that participants displayed during the experiment. Some participants described how they reasoned in their answers to the final question. **Example 1** of such reasoning: "If the game reaches my first decision node and if the payoffs are such that I believe that the computer would not play e if its second decision node is reached, then I play d at my current decision node". This kind of strategic reasoning can be expressed using the following formal notions.

Let us assume that actions are part of the observables, that is, $\Sigma \subseteq P$. The semantics for the actions can be defined appropriately. Let $n_1, \ldots, n_4$ denote the four decision nodes of game 1, with C playing at n_1 and n_3, and P playing at the remaining two nodes n_2 and n_4. We have four belief operators for this game - two for each player. We abbreviate some formulas which describe the payoff structure of the game:

$$\langle d\rangle\langle f\rangle\langle h\rangle((u_C = p_C) \wedge (u_P = p_P)) = \alpha$$
$$\langle d\rangle\langle f\rangle\langle g\rangle((u_C = q_C) \wedge (u_P = q_P)) = \beta$$
$$\langle d\rangle\langle e\rangle((u_C = r_C) \wedge (u_P = r_P)) = \gamma$$
$$\langle c\rangle((u_C = s_C) \wedge (u_P = s_P)) = \delta$$
$$\langle b^-\rangle\langle a\rangle((u_C = t_C) \wedge (u_P = t_P)) = \chi$$
$$\varphi := \alpha \wedge \beta \wedge \gamma \wedge \delta \wedge \chi$$

Let ψ_i denote the conjunction of all the order relations of the rational payoffs for player i given in game 1. A strategy specification describing the strategic reasoning of Example 1 (at the node n_2) is:

$$\eta_P^1 : [(\varphi \wedge \psi_P \wedge \psi_C \wedge \langle b^-\rangle\mathbf{root} \wedge \mathbb{B}_{g1}^{n_2,P}\langle d\rangle\neg e \wedge \mathbb{B}_{g1}^{n_2,P}\langle d\rangle\langle f\rangle g) \mapsto d]^P$$

A BI reasoning at the same node can be formulated as follows:

$$\eta_P^2 : [(\varphi \wedge \psi_P \wedge \psi_C \wedge \langle b^- \rangle \mathbf{root} \wedge \mathbb{B}_{g1}^{n_2,P} \langle d \rangle e \wedge \mathbb{B}_{g1}^{n_2,P} \langle d \rangle \langle f \rangle g) \mapsto c]^P$$

The example above shows how strategic reasoning of participants can be formulated in the proposed framework (which could then be converted to appropriate reasoning rules to build up computational cognitive models). Note that our representations have become quite succinct using the belief operator, compared to the representations we had in [6], because expressions for response strategies are not needed anymore. We leave the details for future work.

4 Modelling in ACT-R

We now provide a brief description of the cognitive architecture at the basis of our computational cognitive model. ACT-R is an integrated theory of cognition as well as a cognitive architecture that many cognitive scientists use [1]. ACT-R consists of modules that link with cognitive functions, for example, vision, motor processing, and declarative processing. Each module maps onto a specific brain region. Furthermore, each module is associated with a buffer and the modules communicate among themselves via these buffers.

The computational cognitive models that we propose are inspired by [6]. We consider a class of models, where each model is based on a set of strategy specifications that can be generated using the logical framework we presented in Section 3. The specifications can represent both backward induction reasoning or forward induction reasoning (in particular, *EFR* reasoning), among others.

Each of the specifications defined in Subsection 3.2 comprises comparisons between relevant payoffs for both the players. For each comparison, a cognitive model has a set of production rules that specify what the model should do. To compare player C's payoffs, say at two leaf nodes, the model first has to find, attend, and encode them in the so-called problem state buffer [1]. For each subsequent payoff, the model performs the following procedure (cf. Figure 5):

- request the visual module to find the payoffs' visual locations;
- direct visual attention to that location; and
- update the problem state (buffer).

The specifications η_P^1 and η_P^2 (see Subsection 3.2) specify what the model should do after encoding the payoffs in the problem state. First, the payoffs need to be compared and the comparison needs to stored. Then the belief operators are dealt with as follows (cf. Figure 5):

- attend visual location of the node depicted by the belief operator; and
- encode the actions and beliefs at the problem state (buffer).

The decisions are made corresponding to the recorded payoffs and the resulting beliefs. An example production rule could be as follows; the model will select and fire this production rule to generate a response:

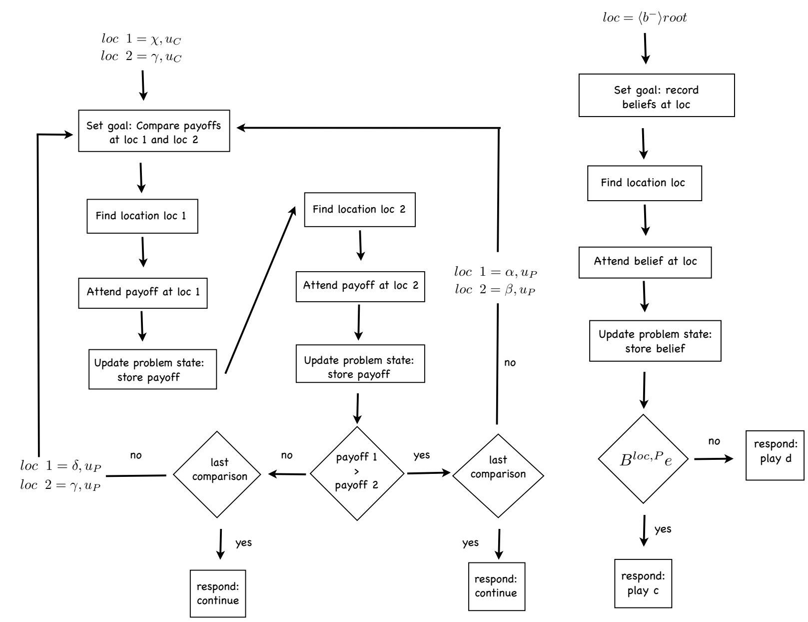

Fig. 5. Flowcharts for reasoning processes as described in Example 1 and BI

IF

Goal is to record Player P's belief at node n	If the current goal is to record Player P's beliefs at node n,
Problem State represents Player P's actions at n, c and d	and the problem state has stored the actions,
$\mathbb{B}^{(P,n)} f$	and belief is f will be played (by C),

THEN

Decision is play d	then request the manual (or motor) module to produce a key press (i.e., play d).

5 Conclusion

In this paper we have continued the line of work started in [6] and proposed another logical language to aid in the construction of computational cognitive models based on the findings of a game-theoretic experiment. We have shown that logic can play a major role in Marr's computational and algorithmic levels of inquiry for cognitive sciences [9]. In future we aim to implement various sets of specifications in separate models, and to simulate repeated game play to study possible learning effects. An advantage of constructing ACT-R models, not only logical formulas, is that quantitative predictions are generated, for example, concerning decision times, locus of attention and activity of brain regions, which can then be tested in further experiments.

References

1. J. Anderson. *How Can Human Mind Occur in the Physical Universe?* OUP, 2007.
2. P. Battigalli. Strategic rationality orderings and the best rationalizability principle. *Games and Economic Behavior*, 13:178–200, 1996.
3. G. Bergwerff, B. Meijering, J. Szymanik, R. Verbrugge, and S. Wierda. Computational and algorithmic models of strategies in turn-based games. In *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, 2014.
4. G. Bonanno. Axiomatic characterization of the AGM theory of belief revision in a temporal logic. *Artificial Intelligence*, 171(2-3):144–160, 2007.
5. S. Ghosh, A. Heifetz, and R. Verbrugge. Do players reason by forward induction in dynamic perfect information games? In R. Ramanujam, editor, *Proceedings of the 15th Conference on Theoretical Aspects of Rationality and Knowledge*, 2015. `http://www.ai.rug.nl/SocialCognition/papers/#conference`.
6. S. Ghosh, B. Meijering, and R. Verbrugge. Strategic reasoning: Building cognitive models from logical formulas. *Journal of Logic, Language and Information*, 23(1):1–29, 2014.
7. S. Ghosh and R. Ramanujam. Strategies in games: A logic-automata study. In N. Bezanishvili and V. Goranko, editors, *Lectures on Logic and Computation - ESSLLI 2010, ESSLLI 2011, Sel. Lecture Notes*, pages 110–159. Springer, 2012.
8. T. Halder, K. Sharma, S. Ghosh, and R. Verbrugge. How do adults reason about their opponent? Typologies of players in turn-taking games. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, 2015. `http://www.ai.rug.nl/SocialCognition/papers/#conference`.
9. D. Marr. *Vision*. Freeman and Company, New York, 1982.
10. R. D. McKelvey and T. R. Palfrey. An experimental study of the centipede game. *Econometrica: Journal of the Econometric Society*, pages 803–836, 1992.
11. B. Meijering, N. A . Taatgen, H. van Rijn, and R. Verbrugge. Modeling inference of mental states: As simple as possible, as complex as necessary. *Interaction Studies*, 15(3):455–477, 2014.
12. B. Meijering, H. Van Rijn, N. A. Taatgen, and R. Verbrugge. I do know what you think I think: Second-order theory of mind in strategic games is not that difficult. In *Proc. 33rd Annual Conf. Cognitive Science Society*, pages 2486–2491, 2011.
13. B. Meijering, H. Van Rijn, N.A. Taatgen, and R. Verbrugge. What eye movements can tell about theory of mind in a strategic game. *PloS ONE*, 7(9):e45961, 2012.
14. R. Nagel and F. F. Tang. Experimental results on the centipede game in normal form: An investigation on learning. *Journal of Mathematical Psychology*, 42(2):356–384, 1998.
15. M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
16. E. Pacuit. Dynamic models of rational deliberation in games. In J. van Benthem, S. Ghosh, and R. Verbrugge, editors, *Models of Strategic Reasoning: Logics, Games and Communities*, volume 8972 of *LNCS-FoLLI*. Springer, New York, NY, 2015.
17. D. Pearce. Rationalizable strategic behaviour and the problem of perfection. *Econometrica*, 52:1029–1050, 1984.
18. A. Perea. *Epistemic Game Theory: Reasoning and Choice*. CUP, 2012.
19. R. Ramanujam and S. Simon. A logical structure for strategies. In *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, volume 3 of *Texts in Logic and Games*, pages 183–208. Amsterdam University Press, 2008.
20. J. van Benthem, S. Ghosh, and F. Liu. Modelling simultaneous games with dynamic logic. *Knowledge, Rationality and Action*, 165:247–268, 2008.