

Preference at First Sight

Chanjuan Liu

School of Computer Science and Technology, Dalian University of Technology

Institute for Logic, Language and Computation, University of Amsterdam

chanjuan.pkucs@gmail.com

We consider decision-making and game scenarios in which an agent is limited by his/her computational ability to foresee all the available moves towards the future – that is, we study scenarios with *short sight*. We focus on how short sight affects the logical properties of decision making in multi-agent settings. We start with *single-agent sequential decision making* (SSDM) processes, modeling them by a new structure of ‘*preference-sight trees*’. Using this model, we first explore the relation between a new natural solution concept of *Sight-Compatible Backward Induction* (SCBI) and the histories produced by classical Backward Induction (BI). In particular, we find necessary and sufficient conditions for the two analyses to be equivalent. Next, we study whether larger sight always contributes to better outcomes. Then we develop a simple logical special-purpose language to formally express some key properties of our preference-sight models. Lastly, we show how short-sight SSDM scenarios call for substantial enrichments of existing fixed-point logics that have been developed for the classical BI solution concept. We also discuss changes in earlier modal logics expressing ‘surface reasoning’ about best actions in the presence of short sight. Our analysis may point the way to logical and computational analysis of more realistic game models.

1 Introduction

There is a growing interest in the logical foundations, computational implementations, and practical applications of *single-agent sequential decision-making* (SSDM) problems [32, 23, 28, 18, 9, 26, 24] in such diverse areas as Artificial Intelligence, Control, Logic, Economics, Mathematics, Politics, Psychology, Philosophy, and Medicine. Making decisions is central to agents’ routine and usually, they need to make multiple decisions over time. Indeed, a current situation is a result of past sequentially linked decisions, each impacted by the preceding choices.

It is quite natural in sequential decision-making scenarios, particularly, in large systems, that agents may have some uncertainties and limitations on their precise view of the environment. The current literature [9, 28] has studied uncertainty which an agent faces in recognizing possible outcomes after taking an action and the probabilities associated with these outcomes, as well as the partial observability of what the actual state is like. In addition to these, a realistic aspect that affects a SSDM process is the short-sightedness of the agent, which blocks a full view of all the available actions. Short sight plays a critical role in such a situation, since, while making a choice, the ability to foresee a variety of alternatives and predict future decision sequences for each of them, may make a significant difference. Nonetheless, such restrictions have not been discussed systematically yet in decision theory or game theory.

In [14], a game-theoretic framework called *games with short sight* was proposed. This framework explicitly models players’s limited foresight in extensive games and calls for a new solution termed as *Sight-Compatible Backward Induction* (SCBI). However, many essential issues related to sight remain unclear, such as: What is the exact role of sight? Will the outcome be better when sight is larger? What is the relation between SCBI and classical *backward induction*(BI)? There are also unexplored issues

pertaining to logical aspects. Which minimal logic is needed for formally characterizing a short-sight framework? Are existing logics for BI still applicable, or can they be extended to fit short-sight scenarios? How different are the logical properties of the game frames for SCBI and for BI? Without such a logical analysis, the framework of [14] does not suffice for disclosing the general features of short sight and the changes it brings about in thinking about decisions and games. Additionally, in multi-player games, short sight has to interact with many other factors, such as agents' mutual knowledge and interactive decisions and moves.

Having said this, we still start by focusing on short sight in single-agent sequential decision-making process. For this, we propose a model of '*preference-sight trees*' (P-S trees). As the term says, a P-S tree combines the agent's *preference* and its *sight*, as both are essential to decision problems [33]. We will study how the two are correlated, and cooperate to act on decision-making processes and their final outcomes.

As a preliminary illustration, consider the connection between larger sight and better outcome. A first impression might be that an agent will always perform better with larger sight. Surprisingly, this is not always true. Sometimes, one can see much further into the future but receive a small payoff, while having one's vision restricted to a limited set of future alternatives yields a better payoff.

Example 1.1 *Alice has to make sequential decisions at two stages (shown in Figure 1). For each stage, she can choose either L or R. Assume that the preference order (from most preferable to least preferable) among the four outcomes is RR, LL, RL, LR. Now consider two cases:*

Case 1. At the start, Alice sees two paths, LR and RL. She chooses R since it initiates RL which is preferable to LR. At the second-stage, Alice then foresees RR and RL. She happily makes the best decision RR.

Case 2. Alice sees more, e.g., LL, LR, and RL, immediately at the first stage. Therefore she thinks that L is a better initial choice than R. Consequently, at the second-stage, she can only choose from LL and LR.

Conclusion: Even though Alice could see more in Case 2, she ultimately obtains a less preferable outcome.




Figure 1: Two-stage decision-making

This example demonstrates some of the crucial features that govern SSDM situations:

1) What an agent can foresee plays a crucial role in the decision-making process, since her sight determines the set of available choices.

2) Sight also updates her preferences over the options, and thereby the outcomes obtained in rational play.

3) Although in Case 2, *Alice* does not get the best result, we can say that, given her sight, she plays optimally in a local sense. In other words, this is a rational plan for her, even though it is not equivalent to the rational outcome of classical decision theory or game theory [29].

In this paper, we address all three challenges, but first we clarify our approach. To focus on sight, we ignore other factors such as the probability of moves by Nature. Also, we model the outcome of a decision as completely determined, or in other words, possible outcomes for each alternative and the probability corresponding to each outcome are encapsulated as a black box.

2 Modeling Single-agent sequential decision-making

We begin by defining a structure called *preference-sight tree* for modelling single-agent sequential decision-making (SSDM) processes. Using this model, we then clarify the role that sight plays by discussing a series of changes it produces in agent's preferences, decision-making procedures and their outcomes.

2.1 Models

There are two kinds of models for decision-making scenarios corresponding to two perspectives. One is an *explicit* model from the perspective of Nature, or an outsider/designer; the other is the *implicit* model from the perspective of the agent involved, or an insider/decider. The former is complete and perfect in the sense that the outsider holds a full view of all the options together with the *objective* quality of these options, and thus can explicitly specify the reward of each situation for the decision-maker. In contrast with this, the latter's views are possibly limited to a near future, especially in large-scale surroundings. Moreover, owing to limited foresight, the agent may also reason mistakenly about the quality of different choices, leading to what we call *subjective* preference.

Both the above perspectives are essential: the former offers a whole picture of the environment, the latter shows the actual play of the decider. In this section, we first introduce an explicit model of *preference trees*. After this, by endowing such trees with the agent's view of the process and his/her subjective preference in this view, we formulate an integrated model of *preference-sight trees* which allows us to model both perspectives together.

2.1.1 Preference trees (P trees)

A preference tree is a decision tree with only two elements: histories and preferences. Each history corresponds to a situation resulting from previous decision actions, and a preference represents the objective quality of each of these situations. To ensure the existence of backward induction solutions, we confine ourselves to finite histories.

Definition 2.1 (*Preference tree*) A **preference tree** is a tuple $T = (H, \succeq)$ where

- H is a non-empty set of finite sequences of actions, called histories.
 - The empty sequence ε is a member of H ;
 - If $(a^k)_{k=1,\dots,K} \in H$ and $L < K$ then $(a^k)_{k=1,\dots,L} \in H$;
- \succeq is a total order over H .

Let A denote the set of all actions. Any history h can be written as a sequence of actions: $(a^k)_{k=1,\dots,n}$, where each $a^k \in A$. If there is no a^{n+1} s.t. $(a^k)_{k=1,\dots,n+1} \in H$, then history $(a^k)_{k=1,\dots,n}$ is a terminal one. The set of terminal histories is denoted Z . The set of actions that are available at h is denoted $A(h) \subseteq A$. For any histories h, h' , if h is a prefix of h' we write $h \triangleleft h'$. The strict part of \succeq is \succ , with $h_1 \succ h_2$ if $h_1 \succeq h_2$ and not $h_2 \succeq h_1$ for any two histories h_1 and h_2 . Accordingly, $h_1 \sim h_2$ iff $h_1 \succeq h_2$ and $h_2 \succeq h_1$.

Several remarks need to be made on the role of preference relations in the above definition:

(1) Instead of defining preference merely over terminal histories, we have defined it over all histories, an idea going back to [19]. Here preference over intermediate histories is necessary for our aim of modelling an agent's decision-making under limited foresight, which usually consists of intermediate histories.

(3) For convenience, we do not strictly differentiate the two main views of preference: qualitative and quantitative. Although we use qualitative order generally, we sometimes switch to numerical payoff when it is advantageous.¹

2.1.2 Preference-sight Trees (P-S trees)

P tree is an explicit model for decision-making scenarios which is independent of an agent. However, for an agent, the tree may appear differently in his/her limited view. [14] proposes the idea of short sight, where the authors use a sight function to denote the set of states that players can actually see at every position in an extensive game. Let us start by adapting their technique to preference trees.

Definition 2.2 Let $T = (H, \succeq)$ be a preference tree. A **sight function** for T is a function $s : H \rightarrow 2^H \setminus \{\emptyset\}$ satisfying $s(h) \subseteq H|_h$ and $|s(h)| < \omega$, where $H|_h$ represents the set of histories extending h . As a special case, $h \in H|_h$.

In words, the function s assigns to each history h a finite subset of all available histories extending h .

The first effect that sight produces is as follows: Given a P tree, for any history h , a sight function always gives us a restricted tree.

Definition 2.3 Let $T = (H, \succeq)$ be a preference tree. Given any history h of T , a **visible tree** T_h of T at h is a tuple (H_h, \succeq_h) , where $H_h = s(h)$, i.e., H_h captures the decider's view of the decision tree; \succeq_h represents the subjective preference over H_h .

A visible tree is actually an implicit model in our earlier terms. H_h also contains a set of terminal histories Z_h , which are those without successors in $s(h)$. Note that typically, the Z_h are non-terminal for T .

Further, the preference order \succeq_h is different from the objective preference \succeq . In fact, the formation of \succeq_h is an update via a bottom-to-top process in terms of an agent's sight. This updating process involves leaving the payoffs of Z_h as the same as their objective payoffs, then updating the payoffs of other histories in H_h backwards, starting from the leaf nodes and proceeding towards the root of the tree.

The reason why we employ such an updating process is that, while the objective payoffs reflect the goodness of these situations, they are not the actual reward that an agent can get if he/she chooses this option. At each decision point, the subjective payoff of one available option is inherited from the best reachable terminal histories of the current visible tree. Therefore, the preference relation \succeq_h in T_h is not always consistent with the preference relation \succeq in T .

This updating process is described by Algorithm 1:

*For convenience, here we use payoffs P to represent rewards.

Fact 2.1 Let $T = (H, \succeq)$ be a P tree. Each visible tree $T_h = (H_h, \succeq_h)$ is a P tree.

Correspondingly, we denote the prefix relation in T_h by \triangleleft_h , and the actions that are available at h by $A_h(h)$.

Finally we proceed to define our model of preference-sight trees. A preference-sight tree allows us not only to represent the outsider's view, i.e., (H, \succeq) , but also to derive a series of implicit models, i.e., (H_h, \succeq_h) , one for each h .

¹There is a debate on whether preference and utilities are the same [18, 7]. Here we adopt the operational understanding of utility and do not distinguish it from preference.

Algorithm 1: Preference updating in visible trees

```

1 PU( $T, h, s$ )
   Input: A P tree  $T = (H, \succeq)$  (or  $T = (H, P)$ ), current history  $h$ , and a sight function  $s$ 
   Output: A visible tree  $T_h = (H_h, \succeq_h)$  or  $(T_h = (H_h, P_h))$ 
2 begin
3    $H \cap s(h) \rightarrow H_h$ ;
4   for any  $z \in Z_h$  /* Keep the payoffs of terminal histories unchanged */ do
5      $P(z) \rightarrow P_h(z)$ ;  $1 \rightarrow flag[z]$ ;
6   while  $flag[h] == 0$  do
7     for any  $h' \in H_h$  do
8       if (for all  $(h'a) \in H_h, flag[(h'a)] == 1$ ) /* If all of its children have been
9         visited, reset its payoff as the highest one among them */
10      then
11         $\max\{P_h(h'a)\} \rightarrow P_h(h')$ ;  $1 \rightarrow flag[h']$ ;
12   Return  $T_h$ ;

```

Definition 2.4 (*Preference-sight tree*) A **preference-sight tree** (*P-S tree*) is a tuple (T, s) , where $T = (H, \succeq)$ is a preference tree and s a sight function for T .

In P-S trees, an agent's sight should satisfy the following properties: First, if an agent can see a given future history, then he/she can also see any intermediate history up to that point. Second, if the agent can see a history two steps forward, then after moving one step ahead, he/she can still see it. These features are formally stated as follows.

Fact 2.2 (*Properties of sight function*) Let (T, s) be a P-S tree. For all $h, h', h'' \in H$, with $h \triangleleft h' \triangleleft h''$, s satisfies :

DC (*Downward-Closed*): if $h'' \in s(h)$, then $h' \in s(h)$.

NF (*Non-Forgetting*): if $h'' \in s(h)$, then $h'' \in s(h')$.

2.2 Solution concepts

Solution concepts are at the center of all choice problems. In what follows, we define two solution concepts for P-S trees, adapted from [30, 14]. After this, we investigate the conditions for their equivalence.

2.2.1 BI history and SCBI history

Backward Induction (BI) is well-known in game theory [30]. The process runs like this. First, one determines the optimal strategy of the player who makes the last move of the game. Using this information, one can then determine the optimal action of the next-to-last moving player. The process continues backwards in this way until all players' actions have been determined in the whole game. Its adaptation to single-agent decision-making process becomes a maximality problem for the agent involved.

In a P-S tree, we say that one history h is \max_{\succeq} in a set of histories $\Gamma \subseteq H$, if $h \in \Gamma$ and for any other history h' in Γ , it holds that $h \succeq h'$, and we write this as $h \in \max_{\succeq} \Gamma$. The strict part for \max_{\succeq} is \max_{\succ} .

Definition 2.5 (BI history) *Let (T, s) be a P-S tree. A history $h^* \in Z$ is a **BI history** of T , iff $h^* \in \max_{\succeq} Z$. Also, we use **BI** to denote the set of BI histories in T .*

A BI history of a P-S tree is a terminal history that is most preferable or equivalently, that has a maximal payoff.

Backward induction precludes short-sight, while in practice it is impossible for an agent to foresee all final outcomes all the time. In [14], a new solution concept was proposed to capture optimal play of short-sighted players: *sight-compatible subgame perfect equilibrium*. The main idea is that at each decision point, the current player chooses a locally optimal move by a local BI analysis within the visible part. Here, we adapt this notion to P-S trees, yielding the *sight-compatible backward induction history*.

Definition 2.6 (SCBI history) *Let (T, s) be a P-S tree. A history $h^* \in Z$ is a **Sight-Compatible Backward Induction history** (SCBI history) of T , iff for each history h with $h \triangleleft h^*$, and the action a following h , i.e., $(ha) \triangleleft h^*$, we have that $\exists z \in \max_{\succeq} Z_h$ such that $(ha) \triangleleft z$. Also, we use **SCBI** to denote the set of SCBI histories in T .*

The difference between SCBI and BI histories is obvious. A BI history is one with highest payoff among the set of terminal histories in the P-S tree, while for a SCBI history every restriction of it should be a local BI history for the visible tree. Thus, BI histories are the BI outcomes for the objective model (H, \succeq) , while SCBI histories are a combination of best responses to all subjective models (H_h, \succeq_h) . Typically it is the case that **SCBI** \neq **BI**.

Example 2.1 *Consider the P-S tree (T, s) in Figure 2, where $s(\varepsilon) = \{L\}$, and $s(L) = \{LR\}$. It is easy to check that **BI** \neq **SCBI**, since **BI** = $\{LL\}$, while **SCBI** = $\{LR\}$.*




Figure 2: **BI** \neq **SCBI**

However, sometimes the two notions can be equivalent.

Example 2.2 *Consider a P-S tree, with T and s shown by Figure 3 (a), and Figure 3 (b) respectively. In (b) the three dotted circles represent $s(\varepsilon)$, $s(L)$ and $s(R)$. For histories L and R , their objective payoffs in (a) are 1 and 2, respectively. However, in T_ε , the subjective payoff of L is updated to 3 and R to 2. Obviously, **BI** = **SCBI** = $\{LL\}$.*




Figure 3: (**BI** = **SCBI**)

2.2.2 Equivalence condition

Then an interesting question on BI and SCBI histories arises: are there conditions under which the two will be equivalent? To get a feeling for this, a first attempt at an answer looks for a condition related to consistency between subjective and objective preferences.

Two histories are said to be ‘preference-sight consistent’ if the subjective preference in each sight-restricted tree is consistent with the objective preference over them:

Definition 2.7 (Preference-sight consistency) Let (T, s) be a P-S tree, and T_h be the visible tree at an arbitrary history h . Then for any two histories h_1, h_2 of T_h , we say (h_1, h_2) satisfies **preference-sight consistency** at h iff

$$h_1 \succeq h_2 \text{ iff } h_1 \succeq_h h_2$$

If for any history $h \in T$, the pair of arbitrary two histories (h_1, h_2) in T_h is *preference-sight consistent* (at h), then we say (T, s) is *preference-sight consistent*.

Is preference-sight consistency an appropriate condition for **BI = SCBI**? We have the following observation:

Fact 2.3 *Preference-sight consistency does not guarantee that **BI = SCBI**.*

Proof 2.4 *Consider Figure 2. Suppose that $s(R)$ contains only one successor. Then it is easy to see that (T, s) is preference-sight consistent. However, **BI** \neq **SCBI**.*

Next, does the other direction hold?

Fact 2.5 *Preference-sight consistency does not follow from **BI = SCBI**.*

Proof 2.6 *The situation in Figure 3 is a counterexample, in which **BI = SCBI** = $\{LL\}$, but (T, s) is not preference-sight consistent, since $R \succ L$ and $L \succ_\varepsilon R$.*

What is the exact condition for **BI = SCBI**? From the failure of preference-sight consistency, we can draw a lesson. In Figure 2, the main reason for (T, s) being inconsistent is that at history L , the branch LL , which in fact forms a BI history, is non-observable to the agent. This tells us that the one with maximal payoff should always be visible. Consider then the example in Figure 3. Here all the options are within agent’s sight, but we notice that although the path LL following L finally turns out to be better than that following R , which makes subjectively $L \succ_\varepsilon R$, the objective payoff of L itself is lower than R . Thus, it fails to imply the consistency between preference and sight.

Based on the above analysis, we now isolate necessary and sufficient conditions for **BI = SCBI**. First, we define an auxiliary property of *sight-reachability*, which intuitively reflects whether each restriction of a history is visible.

Definition 2.8 (*Sight-reachability*) A BI history h^* is **sight-reachable** if, for all $(ha) \triangleleft h^*$, we have $(ha) \in H_h$, where h, h' are histories, and a is an action following h .

Theorem 2.7 (*Equivalence Theorem*) For any P-S tree (T, s) , **SCBI = BI** iff the following conditions are satisfied:

- I). Any history $h^* \in \mathbf{BI}$ is sight-reachable.
- II). Any history $h^* \in \mathbf{BI}$ is locally optimal: For any history $(hh') \triangleleft h^*$, if $(hh') \in Z_h$, then $(hh') \in \max_{\succ} Z_h$ and for any other $(hh'') \in Z_h$, $(hh') \sim (hh'')$ iff $\exists z \in \mathbf{BI}$ such that $(hh'') \triangleleft z$.

Proof 2.8 (\Rightarrow) I). We show that every $h^* \in \mathbf{BI}$ is sight reachable. That is, for all $(hh') \triangleleft h^*$, it holds that $(ha) \in H_h$. By **SCBI = BI**, we know that any history h^* in **BI**, is also in **SCBI**. By Definition 2.6, for each of its prefix h , h_h^* is \max_{\succeq} in Z_h . So h_h^* is in Z_h . In addition, by non-emptiness of Z_h , h_h^* is not an empty sequence. Thus, for all $(ha) \triangleleft h^*$, it holds that $(ha) \in H_h$. So $h^* \in \mathbf{BI}$ is sight-reachable.

To show condition II), take any h^* in **BI**, we have that it is in **SCBI**. Thus, for all $(hh') \triangleleft h^*$, if $(hh') \in Z_h$, then (hh') is \max_{\succeq} in Z_h . Moreover, for any $(hu) \in Z_h$ such that $(hh') \sim (hu)$, we have (hu) is a prefix of a BI history, i.e., $(hu) \in \mathbf{BI}_h$. For suppose not, then (hu) is not a prefix of SCBI history. Then it must be $(hh') \succ (hu)$. Contradict.

(\Leftarrow) Suppose conditions I) and II) are satisfied. It suffices to show (a) “every BI history is SCBI history of T ”, and (b) “every SCBI history is BI history of T ”.

For (a), take any BI history h^* . By I), all BI histories are sight reachable. Further by II), for all $(hh') \triangleleft h^*$, if $(hh') \in Z_h$, then (hh') is \max_{\succeq} in Z_h . This is to say that for each of its prefix h , h_h^* is \max_{\succeq} in Z_h . By definition 2.6, h^* is a SCBI history.

For (b), take any SCBI history h^* . We can show it is a BI history, i.e., h^* is \max_{\succeq} in Z . For suppose not, then there exists a BI history h' such that $h' \succ h^*$. Notice that there must be some history u which is the common prefix of h^* and h' . Since h' is a BI history, by condition I) and II), we know that $h'_u \succ h_u^*$. Then h_u^* is not a prefix of a SCBI history. Thus, h^* is not a SCBI history. Contradiction.

2.2.3 More sight, better outcome?

We have seen earlier on that, SCBI may loss global optimality. The BI history definitely has a maximal payoff, while it might not be the case for SCBI, since each action is chosen with a limited sight. So $\mathbf{BI} \succeq \mathbf{SCBI}$ holds without exception, in the sense that any BI history is no worse than any SCBI history. One might conjecture that more sight always contributes to better outcomes. Yet, the fact below falsifies this.

Fact 2.9 Let T be a P tree. Also, let s_1 and s_2 be two sight functions for T satisfying $s_1(h) \subseteq s_2(h)$ for any history h in T . Take any two SCBI histories z_1 and z_2 of (T, s_1) and (T, s_2) respectively. Then the following three cases are all possible: a) $z_1 \succ z_2$; b) $z_2 \succ z_1$; c) $z_1 \sim z_2$.

Proof 2.10 Case (a) has been shown in Example 1.1. Case (b): Obviously, Figure 2 offers an instance for this. Case (c): The scenario depicted in Figure 3 is an example. \square

In conclusion, full sight guarantees a maximal payoff. However, with short sight, increase of sight does not always improve the outcome. The added sight may bring misleading information, e.g., a branch which is temporarily nicer but actually unpromising, and finally gives rise to an even worse outcome. Still, this does not mean that SCBI is deficient: rather, these observations seem realistic for real agents. These issues will be discussed further in Section 4.

3 A Logical Analysis

After modelling decision-making with short sight by preference tree models, it is instructive to see what a logical language looks like for reasoning about these models, especially the role of sight in a SSDM process. So far, no such logic has been proposed, though logics of game-theoretic structures have been extensively studied – see [22, 20] – while there are a few preliminary logic analyses of sight on its own, [10, 27]. In this section, we design a minimal and natural logical system that supports reasoning about sight in the context of single-agent decision-making processes, characterizing basic properties of preference-sight trees, and formally capturing the results in the previous section.

3.1 Syntax and Semantics

To reason about the key ingredients (i.e., histories, preferences, and sights) of a P-S tree, we take $P^{(T,s)}$ as a set of propositional letters, which at least contains the following ²:

- \overline{h} for each history h .
- $\overline{h_1 \succeq h_2}$ encoding the preference relation of the agent over all histories, and the strict part of which is $\overline{h_1 > h_2}$.
- $\overline{s(h)}$ encoding the sight at each history h in T .

Based on $P^{(T,s)}$, we give a language \mathcal{L} for reasoning about P-S trees. In \mathcal{L} , we have a key dynamic operator $[\!|\varphi]$ for restricting to the worlds satisfying φ , and a universal modality with $A\varphi$ saying that φ is true in every world.

Definition 3.1 (Preference-sight language) Take any set of atomic letters $P^{(T,s)}$. The **preference-sight language** \mathcal{L} is given by the following BNF, where $p \in P^{(T,s)}$:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid [\!|\varphi]\psi \mid A\varphi.$$

We write $\langle\!|\varphi\rangle\varphi$ to abbreviate $\neg[\!|\varphi]\neg\varphi$.

Definition 3.2 (Preference-sight models) For a P-S tree (T, s) , a **preference-sight model** $M^{(T,s)}$ is a tuple $(H, \triangleleft, \mathcal{V})$ where the following holds:

- H is the set of possible worlds, one for each history,
- \triangleleft is the reachability (prefix) relation among worlds,
- $\mathcal{V} : P_T \rightarrow \rho(H)$ is an evaluation function satisfying:

$$(1) \forall h \in H, \mathcal{V}(\overline{h}) = \{h' \mid h' \triangleleft h\}.$$

$$(2) \mathcal{V}(\overline{h_1 \succeq h_2}) = \begin{cases} H, & \text{IF } h_1 \succeq h_2, \\ \emptyset, & \text{Otherwise.} \end{cases}$$

$$(3) \forall h \in H, \mathcal{V}(\overline{s(h)}) = \bigcup_{h' \in s(h)} \mathcal{V}(\overline{h'}).$$

Intuitively, \overline{h} is true at all the worlds leading to h . $\overline{h_1 \succeq h_2}$ is true everywhere if $h_1 \succeq h_2$, and nowhere otherwise. Finally, $\mathcal{V}(\overline{s(h)})$ is a union of the worlds that make the given atom true for at least one element of $s(h)$.

There seems to be nothing striking in this syntax. However, given the special role of atoms, the natural model update differs from the usual one in dynamic-epistemic logic.

Definition 3.3 (Model update) Given a preference-sight model $M^{(T,s)} = (H, \triangleleft, \mathcal{V})$ and a set $X \subset H$, the **updated model** $M_{|X}^{(T,s)}$ produced by the restriction of X is defined as a tuple $(X, \triangleleft \cap X^2, \mathcal{V}_{|X})$, where ³

$$\mathcal{V}_{|X}(p) = \begin{cases} \mathcal{V}_{|X}(\overline{h_1 \succeq h_2}), & \text{IF } p \text{ is of the form } \overline{h_1 \succeq h_2} \\ \mathcal{V}(p) \cap X, & \text{Otherwise} \end{cases}$$

²The idea of defining \overline{h} is motivated by [1], where the authors define an atomic sentence \overline{o} for each leaf in a game tree.

³In this definition, Z_X denotes the terminal histories in X , i.e., the set of histories that have no successors in X .

$$\mathcal{V}_{!X}(\overline{h_1 \geq h_2}) = \begin{cases} X, & \text{IF } \mathcal{V}(\overline{z_1 \geq z_2}) = H, \text{ where} \\ & z_1 \in \max_{\succeq} \{z \in Z_X \mid h_1 \triangleleft z\}, \\ & z_2 \in \max_{\succeq} \{z \in Z_X \mid h_2 \triangleleft z\} \\ \emptyset, & \text{Otherwise} \end{cases}$$

$M_{!X}^{(T,s)}$ is the update of the model $M^{(T,s)}$ restricting the set of states to X , and the valuation function accordingly. But crucially, the valuation for preference atoms in the new model reflects the updating process in the visible tree of Algorithm 1. In the following, we omit superscripts (T, s) .

The semantics for this language is basically standard, [8], so we only mention the truth condition of $[!\phi]\psi$:

Let M be a preference-sight model. For any state h in M ,

$$M, h \models [!\phi]\psi \text{ iff } M, h \models \phi \Rightarrow M_{!X}, h \models \psi.$$

Validity of formulas is defined as usual, cf. [8].

3.2 Main characterization results

Despite its simplicity, \mathcal{L} can express our results in previous sections concerning properties and solutions of P-S trees. We introduce some helpful syntactic abbreviations, and then state our main characterization results.

- $\overline{Z_h} = \bigvee \{ \overline{z} \mid z \in Z_h \}$.
- $\overline{\max_{\succeq} X} = \bigvee \{ \overline{h} \mid h \in X, \text{ and } h \succeq h' \text{ for } \forall h' \in X \}$.
- $\overline{\text{BI}} = \bigvee \{ \overline{z} \mid z \in \text{BI} \}$ ($\overline{\text{BI}}$ holds at T 's BI histories).
- $\overline{\text{SCBI}} = \bigvee \{ \overline{z} \mid z \in \text{SCBI} \}$, that is, the formula $\overline{\text{SCBI}}$ holds at the SCBI histories of T .

Proposition 3.1 *Let (T, s) be a P-S tree and M be a \mathcal{L} -model for it. Then (T, s) is preference-sight consistent iff the following formula is valid in M :*

$$\bigwedge_h \bigwedge_{h_1 \in H_h} \bigwedge_{h_2 \in H_h} ((\overline{h_1 \geq h_2} \rightarrow [!s(h)]\overline{h_1 \geq h_2}) \wedge (\langle !s(h) \rangle \overline{h_1 \geq h_2} \rightarrow \overline{h_1 \geq h_2}))$$

Lemma 3.2 *For any P-S tree (T, s) and model M for it, a BI history h^* is sight-reachable if and only if the following formula holds in M :*

$$(\text{SR}) : \bigwedge_h \bigwedge_{a \in A(h)} (A(\overline{(ha)} \rightarrow \overline{h^*}) \rightarrow (A(\overline{(ha)} \rightarrow \overline{s(h)}))).$$

Proof 3.3 (\Rightarrow) *Suppose that BI history h^* is sight-reachable. By Definition 2.8, we have that, for all $(ha) \triangleleft h^*$, it holds that $(ha) \in s(h)$, where h, h' are histories, and a is an action following h . More formally, $(ha) \triangleleft h^*$ can be defined by the formula $A(\overline{(ha)} \rightarrow \overline{h^*})$ in the sense that, in T , for all h and $a \in A(h)$, $(ha) \triangleleft h^*$ iff $M \models A(\overline{(ha)} \rightarrow \overline{h^*})$. And similarly $(ha) \in s(h)$ is defined by $A(\overline{(ha)} \rightarrow \overline{s(h)})$. Thus if a BI history h^* is sight-reachable, then $M \models \bigwedge_h \bigwedge_{a \in A(h)} (A(\overline{(ha)} \rightarrow \overline{h^*}) \rightarrow (A(\overline{(ha)} \rightarrow \overline{s(h)})))$. The other direction can be proved in a similar way. \square*

Lemma 3.4 Let (T, s) be a P-S tree and M be a \mathcal{L} -model for it. A BI history h^* is locally optimal iff the following formula is valid in M :

$$(LO) : \quad \left(\bigwedge_h \bigwedge_{(hh') \in Z_h} (A(\overline{(hh')} \rightarrow \overline{h^*}) \rightarrow (A(\overline{(hh')} \rightarrow \overline{\max_{\succeq} Z_h}) \wedge \bigwedge_{(hh'') \in Z_h} ((hh') \sim (hh'') \leftrightarrow \bigvee_{z \in \mathbf{BI}} (A(\overline{(hh'')} \rightarrow \overline{z})))) \right)$$

Proof 3.5 (\Leftarrow) Suppose BI history h^* is locally optimal. Then for $(hh') \triangleleft h^*$, if $(hh') \in Z_h$, we have (hh') is \max_{\succeq} in Z_h . And for any (hh'') , $(hh'') \sim (hh')$ iff $\exists z \in \mathbf{BI}$ s.t. $(hh'') \triangleleft z$. Similar with the above proposition, $A(\overline{(hh')} \rightarrow \overline{h^*})$ captures that $(hh') \triangleleft h^*$. And $A(\overline{(hh'')} \rightarrow \overline{z})$ shows that $(hh'') \triangleleft z$. Finally, $(\overline{(hh')} \rightarrow \overline{\max_{\succeq} Z_h})$ demonstrates that (hh') is \max_{\succeq} in Z_h . Direction (\Rightarrow) uses a similar check.

Proposition 3.6 (\mathcal{L} -characterization of equivalence) Let (T, s) be a preference-sight tree and M a model for it. Then the following formula is valid in M :

$$\models (A(\overline{\mathbf{BI}} \leftrightarrow \overline{\mathbf{SCBI}})) \leftrightarrow \bigwedge_{h^* \in Z} ((A(\overline{h^*} \rightarrow \overline{\mathbf{BI}})) \rightarrow (SR \wedge LO))$$

Proof 3.7 Direction (\Rightarrow). We need to prove the following:

- 1) $(A(\overline{\mathbf{BI}} \leftrightarrow \overline{\mathbf{SCBI}})) \rightarrow \bigwedge_{h^* \in Z} (A(\overline{h^*} \rightarrow \overline{\mathbf{BI}}) \rightarrow SR)$.
- 2) $(A(\overline{\mathbf{BI}} \leftrightarrow \overline{\mathbf{SCBI}})) \rightarrow \bigwedge_{h^* \in Z} (A(\overline{h^*} \rightarrow \overline{\mathbf{BI}}) \rightarrow LO)$.

For 1). It is equivalent to prove that, for any $h^* \in Z$, $(\overline{\mathbf{BI}} \leftrightarrow \overline{\mathbf{SCBI}}) \wedge (A(\overline{h^*} \rightarrow \overline{\mathbf{BI}})) \rightarrow SR$. Suppose $\neg(SR)$. Then $\exists(ha) \triangleleft h^*$, and $(ha) \notin T_h$, and so, at h , the branch leading to h^* is not visible in T_h . Thus, the BI history in T_h could not be a branch leading to h^* . By the definition **SCBI**, it follows that $h^* \notin \mathbf{SCBI}$. However, by $\overline{h^*} \rightarrow \overline{\mathbf{BI}}$ we know that h^* is a BI history. This contradicts $\overline{\mathbf{BI}} \leftrightarrow \overline{\mathbf{SCBI}}$.

2) can be proved in a similar style.

Direction (\Leftarrow). Suppose that $\neg(A(\overline{\mathbf{BI}} \leftrightarrow \overline{\mathbf{SCBI}}))$. Then

- (a): $\exists z^* \in \mathbf{BI}$ and $z^* \notin \mathbf{SCBI}$, or
- (b): $\exists z^* \in \mathbf{SCBI}$ and $z^* \notin \mathbf{BI}$.

If (a), then, by the antecedent, we have that: $\forall(ha) \triangleleft z^*$, $(ha) \in H_h$. Also, $\forall(hh') \in Z_h$ and $(hh') \triangleleft h^*$, it holds that $(hh') \in \max_{\succeq} Z_h$. Then it directly follows that z^* is a SCBI history. Contradiction.

If (b), then take any $z \in \mathbf{BI}$, which shares a prefix u with z^* , i.e., $u \triangleleft z$ and $u \triangleleft z^*$. By the antecedent, we have $z_u \in \max_{\succeq} Z_h$. Since $z^* \notin \mathbf{BI}$, it follows that $z_u > z_u^*$. Then $z^* \notin \mathbf{SCBI}$. Once more, we have a contradiction. \square

3.3 Valid principles

The operator $[\!|\varphi|]$ makes \mathcal{L} a PAL-like language. However, the special model-update makes it different from standard PAL [11]. This suggests a close look at what is and what is not valid in preference-sight models.

First, some axioms in standard PAL do not hold in preference-sight models. For example, the **!ATOM** axiom, $[\!|\varphi|]p \leftrightarrow (\varphi \rightarrow p)$, is not valid when it is of the form below.

Proposition 3.8 *The following is not valid in preference-sight models, where h, h_1, h_2 represent arbitrary histories.*

$$\text{!Sight-Preference : } \quad [!s(\bar{h})]\bar{h}_1 \geq \bar{h}_2 \leftrightarrow (\overline{s(\bar{h})} \rightarrow \overline{h_1 \geq h_2}).$$

Proof 3.9 *For a counterexample, consider the tree T in Figure ???. It is easy to see that in the model M for T , $M \models [!s(\bar{\varepsilon})]\bar{h}_1 \geq \bar{h}_2$ and $M \not\models \overline{s(\bar{\varepsilon})} \rightarrow \overline{h_1 \geq h_2}$, since there exists a state ε such that $M, \varepsilon \models \overline{s(\bar{\varepsilon})}$ and $M, \varepsilon \not\models \overline{h_1 \geq h_2}$.*

This proposition says that subjective preference in visible trees is not necessarily consistent with objective preference.

Now let us see some interesting valid principles and their intuitive interpretations.

Lemma 3.10 *The formulas shown in Table 1 are valid, where h, h_1, h_2 , and h_3 are arbitrary histories.*

$Taut$	all propositional tautologies
T_{\geq}	$\overline{h \geq h}$
4_{\geq}	$\overline{h_1 \geq h_2 \wedge h_2 \geq h_3} \rightarrow \overline{h_1 \geq h_3}$
to_{\geq}	$\overline{h_1 \geq h_2 \vee h_1 \geq h_2}$
T_s	$\bar{h} \rightarrow \overline{s(\bar{h})}$
TM	$\bigwedge_{z \in Z} \bigwedge (A(\bar{z} \rightarrow \bar{h}) \rightarrow A(\bar{h} \rightarrow \bar{z}))$
DC	$\bigwedge_{h_3 \triangleleft h_2 \triangleleft h_1} \bigwedge_{h_3 \triangleleft h_1} (A(\bar{h}_3 \rightarrow \overline{s(\bar{h}_1)}) \rightarrow A(\bar{h}_2 \rightarrow \overline{s(\bar{h}_1)}))$
NF	$\bigwedge_{h_3 \triangleleft h_2 \triangleleft h_1} \bigwedge_{h_3 \triangleleft h_1} (A(\bar{h}_3 \rightarrow \overline{s(\bar{h}_1)}) \rightarrow A(\bar{h}_3 \rightarrow \overline{s(\bar{h}_2)}))$
$!_{ATOM, SP}$	$[! \varphi] p \leftrightarrow (\varphi \rightarrow p)$ (excluding the schema !Sight-Preference)
$!NEG$	$[! \varphi] \neg \psi \leftrightarrow (\varphi \rightarrow \neg [! \varphi] \psi)$
$!CON$	$[! \varphi] (\psi \wedge \chi) \leftrightarrow ([! \varphi] \psi \wedge [! \varphi] \chi)$
$!COM$	$[! \varphi] [! \psi] \chi \leftrightarrow [! \varphi \wedge [! \psi] \chi]$
$Dual$	$[! \varphi] \psi \leftrightarrow \neg (! \varphi) \neg \psi$

Table 1: Valid principles of L

Proof 3.11 *We only prove some cases, proofs for the others are trivial or standard.*

For T_s . Take any state u with $M, u \models \bar{h}$. Then $u \in \mathcal{V}(\bar{h})$. As the sight function is reflexive, i.e., $h \in s(h)$, it holds that $\mathcal{V}(\bar{h}) \subseteq \mathcal{V}(\overline{s(\bar{h})})$. So $u \in \mathcal{V}(\overline{s(\bar{h})})$. Thus, $M, u \models \overline{s(\bar{h})}$.

For TM . Take any state u , any history h and any $z \in Z$, and suppose $M, u \models A(\bar{z} \rightarrow \bar{h})$. Then for any u' , $u' \in \mathcal{V}(\bar{z})$ implies that $u' \in \mathcal{V}(\bar{h})$. Thus, $\mathcal{V}(\bar{z}) \subseteq \mathcal{V}(\bar{h})$. It follows that $z \in \mathcal{V}(\bar{h})$. Given that z is terminal, by the definition of $\mathcal{V}(\bar{h})$, it must be that $h = z$. Thus, $M, u \models A(\bar{h} \rightarrow \bar{z})$.

For DC . Take any state u , suppose for some $h_1 \triangleleft h_2 \triangleleft h_3$, $M, u \models A(\bar{h}_3 \rightarrow \overline{s(\bar{h}_1)})$. Then we know $\mathcal{V}(\bar{h}_3) \subseteq \mathcal{V}(\overline{s(\bar{h}_1)})$. It follows that $h_3 \in s(\bar{h}_1)$. As the sight function is downward closed, we have $h_2 \in s(\bar{h}_1)$. Thus, $M, u \models A(\bar{h}_2 \rightarrow \overline{s(\bar{h}_1)})$.

For $!_{ATOM, SP}$. Take any state u , and let $M, u \models [! \varphi] p$ where φ is not of the form $\overline{s(\bar{h})}$ and p is not of the form $\overline{h_1 \geq h_2}$. It holds that $M, u \models \varphi$ implies that $M, u \models p$. By Definition 3.3, $M, u \models [! \varphi] p$ iff $M, u \models \varphi \rightarrow p$. Therefore, $M, u \models \varphi$ implies $M, u \models p$. Equivalently, then, $M, u \models \varphi \rightarrow p$. \square

Interpretation of valid principles. Each of these axioms has some intuitive appeal. T_{\geq} , 4 and to_{\geq} show the *reflexivity*, *transitivity* and *totality* of the preference relation, respectively. Likewise, T_s says that sight is *reflexive*. DC characterizes the (*downward-closure*) property of sight. NF encodes the *non-forgetting* property of sight. TM guarantees that terminal histories of the P-S tree are actually *terminal*. One further interesting point is that there is no correspondence of TM for terminal histories of visible trees.

Fact 3.12 *The following formula is not valid in preference-sight models:*

$$\bigwedge_u \bigwedge_{\bar{z} \in Z_u} \bigwedge_h (A(\bar{z} \rightarrow \bar{h}) \rightarrow A(\bar{h} \rightarrow \bar{z})).$$

Other validities in the table are axioms for standard PAL. We postpone the study of a complete axiomatization of the logic L until future work.

To conclude this section, in \mathcal{L} , the ingredients including histories, preferences and sights are encoded as primitive propositions. Various earlier phenomena in P-S trees can thus be captured in a simple, direct and intuitive manner. This special-purpose logic, as we will see soon, is model-dependent, but it can also be formulated generically.

4 Background in game logics

In this section, we relate our logic L to existing logics for classical game theory, showing how ideas can be combined where useful. Since so far we have been working with BI and SCBI histories, we first define strategies for P-S trees: A *strategy* for a P-S tree (T, s) is a function $\sigma : H \rightarrow A$ such that $\sigma(h) \in A(h)$. That is, σ assigns each history h an action that follows h . In particular, for a visible tree T_h , a ‘local strategy’ σ_h is a restriction of σ to T_h , such that $\sigma_h(h') = \sigma(h')$ for any $h' \in T_h$.

4.1 Generic formulation of \mathcal{L}

In applied logic for structure analysis, there exist two extremes, viz. model-dependent ‘local languages’ and ‘generic languages’ that work across models. For a generic logic, a definition of a property π is a formula φ such that for all models M , M has property π iff $M \models \varphi$. For a local language, such a formula can depend on a given model M : there exists a formula φ_M which depends on M , such that any model M has the property π iff $M \models \varphi_M$. However, in this case, the defining formula can be trivial. For example, one might define φ_M simply as follows.

$$\varphi_M = \begin{cases} \top, & \text{if } M \text{ satisfies } \pi \\ \perp, & \text{Otherwise} \end{cases}$$

In this subsection, using a well-known *Rationality* property as an example, we discuss how model-dependent our earlier language L is, and then show how it can be formulated in a generic way. We first recall the results on classical BI. Given that we have been dealing with single-agent cases until now, in this Section, we will adapt the results from the literature on multi-player games to the single-player case.

The BI strategy [2, 3] is the largest subrelation σ of the total *move* relation that has at least one successor at each node, while satisfying the rationality (RAT) property:

RAT No alternative move for the player yields an outcome via further play with σ that is strictly better than all the outcomes resulting from starting at the current move and then playing σ all the way down the tree.

As argued in [2, 3], this rationality assumption is a confluence property for action and preference:

$$\mathbf{CF} \quad \forall x \forall y (x \sigma y \rightarrow \forall z (x \text{ move } z \rightarrow \exists u (\text{end}(u) \wedge y \sigma^* u \wedge \forall v ((\text{end}(v) \wedge z \sigma^* v) \rightarrow u \geq v))))$$

We can observe that there is also a corresponding rationality property for the local BI strategies that constitute an SCBI, which should however now express a confluence property for action, preference and sight. Specifically, for a P-S tree, each local BI strategy for the visible tree T_h at h is the largest subrelation σ_h of the total *move* relation in T_h , satisfying 1) σ_h has at least one successor at each $h' \in T_h$, and 2) the following rationality property holds:

RATS In the visible tree, there is one outcome obtained by playing σ_h from the start to the end, that is no worse than all the outcomes yielded from any alternative first move followed by further play with σ_h .

This confluence property involving sight is expressible as follows in our language \mathcal{L} :

Proposition 4.1 *Let (T, s) be a P-S tree, and let M be any model for it. M satisfies RATS iff M validates the following \mathcal{L} -formula, where σ_h is the BI strategy for visible tree at h and where $(h(\sigma_h)^k)$ stands for the history reached from h after executing σ_h for k times.*

$$\mathbf{CFS}_M \quad \bigwedge \bigvee_{h} \bigvee_{z \in Z_h} \bigvee_{k=l(z)} \left(A(\overline{(h(\sigma_h)^k}) \leftrightarrow \bar{z}) \right) \rightarrow \left(\bigwedge_{a' \in A_h(h)} \bigwedge_{z' \in Z_h} \bigwedge_{m=l(z')-l(ha')} \left(A(\overline{(ha'(\sigma_h)^m}) \leftrightarrow \bar{z}') \right) \rightarrow \bar{z} \geq \bar{z}' \right),$$

Proof 4.2 *We first claim that at state $h \in H$, for any terminal history $z \in Z_h$, and $h' \in H_h$, $A(\overline{h'} \leftrightarrow \bar{z})$ implies that $h' = z$. This is straightforward since $A(\overline{h'} \leftrightarrow \bar{z})$ demonstrates that prefixes of h' are the same with those of z , which means that $h' = z$. Then $M \models \mathbf{CFS}_M$ says that there is a terminal history z_h following h by playing a local BI strategy σ_h , such that $z \succeq z'$ for any other $z' \in Z_h$ which follows an alternative first move $d' \in A_h(h)$ via further play of σ_h . Therefore, we know that M satisfies **RATS**. \square*

However, compared with the generic logic in [4, 3, 2], the given definition in our logic is local. It is obvious that **CF**, the formula defining the property **RAT**, is insensitive to models – while our **CFS** _{M} relies on a given model for its ranges of big disjunctions and conjunctions, and in its model-dependent notations like $s(h)$ and $\overline{h_1} \geq \overline{h_2}$. Still, it is also clearly true that our definition is not as trivial as the earlier local trick. Therefore, our logic \mathcal{L} seems somewhere between the two extremes of locality and genericity. This feeling can be made precise by moving to a closely related truly generic first-order logic.

The relevant modified formula involves some natural auxiliary predicates. $x \triangleleft y$ says that x is a prefix of y ; $x \triangleleft y$ means that x can see y . Corresponding to the BI relation σ , $y \sigma(x)z$ says that from y , z is a local backward induction move in the visible tree at x ; σ^k describes σ being composed for k times with $k \in \mathbb{N}^4$; *move* and \geq are still the move relation and preference relation, respectively, of the game.

Proposition 4.3 *Any model M satisfies RATS iff it validates the following formula.*

$$\mathbf{CFS}(FO): \quad \forall x \{ (\exists y (x \triangleleft y)) \rightarrow \forall u [(x \sigma(x)u) \rightarrow \forall t ((x \text{ move } t \wedge x \triangleleft t) \rightarrow \exists z ((x \triangleleft z \wedge \neg \exists z' (z \triangleleft z' \wedge x \triangleleft z')) \wedge \exists k (u(\sigma(x))^k z)) \wedge \forall v ((x \triangleleft v \wedge \neg \exists v' (v \triangleleft v' \wedge v \triangleleft v')) \wedge \exists l (t(\sigma(x))^l v)) \rightarrow \wedge z \geq v)) \} \}.$$

⁴Here $x \sigma^k y$ is the abbreviation of $\exists y_1 \exists y_2 \dots \exists y_k (x \sigma y_1 \wedge y_1 \sigma y_2 \wedge \dots \wedge y_{k-1} \sigma y_k \wedge (y_k = y))$.

Proof 4.4 *It is easy to show that $M \models \mathbf{CFS}(FO)$ iff $M \models \mathbf{CFS}_M$. \square*

In summary, incorporating basic elements of P-S trees directly into first-order syntax makes L intuitive and natural.

Even so, other logics exist for dealing with further aspects of game trees and solution procedures, and we will discuss a few examples in what follows with a view to how they behave in the presence of sight.

4.2 Solution procedures and fixed-point logics

Recursive solution procedures naturally correspond to definitions in existing fixed-point logics, such as the widely used system LFP(FO). An LFP(FO) formula mirroring the recursive nature of BI is constructed in [4, 6] to define the classical BI relation, based on the above property **RAT**. Now, we have shown that sight-restricted SCBI, too, is a recursive game solution procedure. Can LFP(FO) be used to define SCBI as well – and if so, how?

The answer is yes, but we need an extension. Rather than a binary relation bi as in [4, 6], characterizing SCBI needs a *ternary* relation. First, we define the local BI relation in visible trees, which will be denoted by bi_{sight} . For any states x, y, z , $bi_{sight}(x, y, z)$ means that in the visible tree at x , the local BI strategy is bi_{sight} , which chooses z when the current state is y . It is then obvious that bi_{sight} should satisfy the following simple first-order definable property, requiring the relevant states to be visible and reachable:

$$bi_{sight}(x, y, z) \rightarrow see(x, y) \wedge see(x, z) \wedge move(y, z).$$

The intuition of $bi_{sight}(x, y, z)$ is then captured as follows:

$$\begin{aligned} & \forall x \forall y \forall z (bi_{sight}(x, y, z) \rightarrow \forall t ((see(x, t) \wedge move(y, t)) \\ & \rightarrow (\exists u (end_{sight}(x, u) \wedge bi_{sight}^*(x, z, u) \wedge \forall v ((end_{sight}(x, v) \wedge \\ & bi_{sight}^*(x, t, v)) \rightarrow u \geq v)))))). \end{aligned}$$

Notice that all occurrences of bi_{sight} in the above formulas are still syntactically positive. This allows us to define local BI strategy bi_{sight} with LFP(FO).

Proposition 4.5 *The strategy bi_{sight} can be defined as the relation R in the following LFP(FO) formula.*

$$\begin{aligned} & \nu R, xyz \bullet \forall x \forall y \forall z (R(x, y, z) \rightarrow \forall t ((see(x, t) \wedge move(y, t)) \\ & \rightarrow (\exists u (end_{sight}(x, u) \wedge R^*(x, z, u) \wedge \forall v ((end_{sight}(x, v) \wedge \\ & R^*(x, t, v)) \rightarrow u \geq v)))))). \end{aligned}$$

It can be proved formally that bi_{sight} is a greatest-fixed-point of the above formula. Based on bi_{sight} , we now proceed to show that the SCBI relation is LFP(FO) definable.

Corollary 4.6 *The SCBI relation $schi$ for a P-S tree can be represented in the following formula:*

$$\forall x \forall y (schi(x, y) \leftrightarrow bi_{sight}(x, x, y)).$$

As in the original classical case, this LFP(FO) definability of $schi$ exposes an intersection between the logical foundation of computation and the recursive nature of sight-compatible backward induction solutions for P-S trees.

4.3 Modal surface logic of best action

In contrast with detailed formalism of solutions with LFP(FO), there is the modal surface logic of [5], which enables direct and natural reasoning about best actions without considering the underlying details of recursive computation. First of all, we list its modalities for classical BI. $[bi]$ and $[BI]$ encode the BI move and BI paths respectively. $[best]\varphi$ says that φ is true in some successor of the current node that can be reached in one step via the bi move.

$$\begin{aligned}
M, h &\models \text{end} \text{ iff } h \in Z. \\
M, h &\models [\text{move}]\varphi \text{ iff } \forall h' = (ha) \text{ with } a \in A(h), M, h' \models \varphi. \\
M, h &\models [\text{best}]\varphi \text{ iff for all } h' \text{ with } h' \in bi(h), M, h' \models \varphi. \\
M, h &\models [bi]\varphi \text{ iff for all } h' \text{ with } h' \in bi(h), M, h' \models \varphi. \\
M, h &\models [bi^*]\varphi \text{ iff } M, u \models \varphi \text{ for all } u \text{ with } u \in (bi)^*(h). \\
M, h &\models [BI]\varphi \text{ iff for all } z \text{ with } z \in \mathbf{BI}, M, z \models \varphi.
\end{aligned}$$

The above logic is still applicable in our setting, but it requires substantial extension for sight-related concepts. In accordance with $[bi]$ and $[BI]$, we use $[scbi]$ and $[SCBI]$ as operators for the SCBI strategy and SCBI path, respectively. For the local BI strategy and path in visible trees, the modalities are $[bi_{sight}]$ and $[BI_{sight}]$. Moreover, recall that $M_{!s(h)}$ is the updated model obtained in the way of Definition 3.3.

$$\begin{aligned}
M, h &\models [scbi]\varphi \text{ iff for all } h' \text{ with } h' \in scbi(h), M, h' \models \varphi. \\
M, h &\models [SCBI]\varphi \text{ iff for all } h' \text{ with } z \in \mathbf{SCBI}, M, z \models \varphi. \\
M, h &\models [!sight]\varphi \text{ iff } M_{!s(h)}, h \models \varphi. \\
M_{!s(h)}, u &\models \text{end}_{sight} \text{ iff } u \in Z_h. \\
M_{!s(h)}, u &\models [\text{move}_{sight}]\varphi \text{ iff for } \forall u' = (ua) \text{ with } a \in A_h(u), \\
&\quad M_{!s(h)}, u' \models \varphi. \\
M_{!s(h)}, u &\models [\text{best}_{sight}]\varphi \text{ iff } M, u' \models \varphi \text{ for } \forall u' \in bi_h(u). \\
M_{!s(h)}, u &\models [bi_{sight}]\varphi \text{ iff } M, u' \models \varphi \text{ for } \forall u' \in bi_h(u). \\
M_{!s(h)}, u &\models [(bi_{sight})^*]\varphi \text{ iff } M_{!s(h)}, u' \models \varphi \text{ for all } u', \\
&\quad \text{such that } u' \in (bi_h)^*(u). \\
M, h &\models [BI_{sight}]\varphi \text{ iff for all } z \text{ with } z \in \mathbf{BI}_h, M, z \models \varphi.
\end{aligned}$$

We give a few illustrations of new issues that arise now.

Capturing the SCBI strategy For a start, we are now able to characterize the SCBI strategy, in a similar vein as the frame correspondence for the classical BI strategy in [5].

Proposition 4.7 *The BI strategy is the unique relation bi satisfying this modal axiom for all propositions p :*

$$(\langle bi^* \rangle(\text{end} \wedge p)) \rightarrow ([\text{move}][\sigma^*](\text{end} \wedge \langle \leq \rangle p))$$

Along the same lines, we can express the SCBI strategy in P-S trees based on the idea that each $scbi$ move coincides with a local BI move within the current visible tree.

Proposition 4.8 *The SCBI strategy is the relation $scbi$ satisfying the following axioms for all propositions p :*

$$\begin{aligned}
(1) \quad &\langle scbi \rangle p \leftrightarrow [!sight]\langle bi_{sight} \rangle p. \\
(2) \quad &[!sight](\langle (bi_{sight})^* \rangle(\text{end}_{sight} \wedge p) \rightarrow \\
&[\text{move}_{sight}]\langle (bi_{sight})^* \rangle(\text{end}_{sight} \wedge \langle \leq \rangle p)).
\end{aligned}$$

Best action and preference-consistency Turning to properties of frames for the extended modal logic of best action with sight, there are interesting differences when comparing SCBI and classical BI. To see this, we employ operators $\langle best \rangle$, $\langle best_{sight} \rangle$, $\langle bi^* \rangle$, $\langle scbi^* \rangle$ and $\langle bi_{sight} \rangle^*$. Now we can make some interesting comparisons.

Proposition 4.9 *For classical backward induction, the axiom $\langle best \rangle \langle bi^* \rangle \varphi \leftrightarrow \langle bi^* \rangle \varphi$ holds.*

However, the new frames do not have the corresponding axiom for the SCBI strategy, since the actions it recommends are not necessarily the actual best actions according to BI. Even in visible trees, this is also not true.

Proposition 4.10 *The following formulas are not valid:*

- (a) $\langle best \rangle \langle scbi^* \rangle \varphi \leftrightarrow \langle scbi^* \rangle \varphi$.
- (b) $[\!|sight|] (\langle best \rangle \langle bi_{sight} \rangle^* \varphi \leftrightarrow \langle bi_{sight} \rangle^* \varphi)$.

Nevertheless, there is a certain coherence between the local BI strategy and local best actions returned by it.

Proposition 4.11 *The following formula is valid: $[\!|sight|] (\langle best_{sight} \rangle \langle bi_{sight} \rangle^* \varphi \leftrightarrow \langle bi_{sight} \rangle^* \varphi)$.*

As for the preference relation, SCBI has a property that classical BI lacks: local BI moves never conflict with the preferences in submodels. In other words, within a visible tree, the initial move determined by the local BI strategy is more preferable for the agent than any other first move.

Proposition 4.12 *For SCBI, it holds that $[\!|sight|] (\langle best_{sight} \rangle \varphi \rightarrow [move_{sight}] \langle \leq \rangle \varphi)$.*

For BI, although it returns a final optimal path, there is no guarantee that its intermediate histories be preferable.

Proposition 4.13 *For BI, the following does not hold: $\langle best \rangle \varphi \rightarrow [move] \langle \leq \rangle \varphi$.*

Path terminality and optimality Using a similar style of modal analysis, we can make the following observations concerning the obvious operators $[BI]$, $[SCBI]$ and $[BI]_{sight}$.

Proposition 4.14 *We have the following three facts:*

- (a) *The formula $[BI] \varphi \rightarrow [BI][BI] \varphi$ is valid.*
- (b) *For SCBI, the following formula does not hold: $[BI]_{sight} \varphi \rightarrow [BI]_{sight}[BI]_{sight} \varphi$.*
- (c) *The formula $[SCBI] \varphi \rightarrow [SCBI][SCBI] \varphi$ is valid.*

Here (a) says that from a BI outcome only a terminal history can be reached; (b) shows that the local BI history may not be a terminal history of the whole tree, and (c) says the SCBI history for the whole tree is always terminal.

Another phenomenon regarding these operators is the local optimality of SCBI at the cost of being more realistic than BI. We have mentioned this point already in Section 2.2.4: now we can present a precise formal version.

Proposition 4.15 *Let σ be any strategy profile,*

- (a). *For BI, the following is valid: $\langle BI \rangle \varphi \rightarrow [\sigma] \langle \leq \rangle \varphi$.*
- (b). *The following does not hold: $\langle SCBI \rangle \varphi \rightarrow [\sigma] \langle \leq \rangle \varphi$.*
- (c). *For SCBI, it holds that $[\!|sight|] (\langle BI_{sight} \rangle \varphi \rightarrow [\sigma_{sight}] \langle \leq_{sight} \rangle \varphi)$.*

Here (a) shows the global optimality of the BI path. (b) and (c) together say the SCBI path is not globally optimal, but each move on this path leads to a locally optimal path.

Altogether, this section has shown the broad logical foundations of our framework, embedding our local language in existing broader generic formalisms, but also enriching and extending these frameworks with aspects of short sight.

5 Toward Multi-player games

While our models and results are about single-agent sequential decision-making processes, we believe they are applicable well beyond that. They can be naturally extended to multi-player extensive game-scenarios with short sight. For such a game model, we can build on [14], which makes an assumption that the current player only knows his own sight, and that he believes other players can see as much as he can see and will play according to this belief. That is, this model precludes more complex forms of interactive knowledge and reasoning. But using this same assumption, our model in this paper can be extended to multi-player cases directly. The only thing we have to do is add agent-labeling to SSDM: even though players can change with time, everything including sight, preference, and actions can be modeled from the current player's perspective.

We will not state any results for the extended multi-player model since they are quite similar to what we have shown already. The case where we drop the above assumption and allow a more free modeling of players' mutual knowledge and beliefs about sight and preference would be more interesting. We will leave this for future work.

6 Discussion and Conclusion

Though motivated by single-agent decision-making process, we have gone towards a much more general goal. In the process, our analysis significantly adds to current connections between logic, computation, and game solutions.

In many recent game-theoretic papers centering on *bounded rationality*, a model has been used of *games with awareness*, [15, 17, 21, 13, 16]. This approach generalizes the classical representation of extensive games by modeling players who may not be aware of all the paths. While [14] shows that games with short sight are a well-behaved subclass of games with awareness, there exists a fundamental difference in focus. Players in the latter approach may be unaware of some branches but they can always see some terminal histories, while in the former, players' sight may only include intermediate histories, ruling out all terminal ones. Moreover, we have shown how short-sight games allow for a natural co-existence of two views of a game, that of insiders and that of outsiders. Having said this, it is clearly an interesting issue to see if our approach in this paper can be extended to cover awareness.

Another obvious interface for our logics are heuristic evaluation approaches for intermediate nodes used by the AI community for computational game-solving, [25, 12, 33]. This, too, is a connection that deserves further exploration.

There are many additional topics to pursue. We already mentioned multi-player scenarios with non-trivial interactive reasoning about other agents' preferences, sights, and strategies. This has also been identified as a key task for epistemic game theory, [31].

Acknowledgments

I thank Fenrong Liu for our fruitful collaboration on earlier versions of this paper. Paolo Turrini provided crucial insights on short-sight games and their connections with games and computation, which we are partly exploring together. Sonja Smets provided helpful comments overall. But especially, I thank Johan van Benthem for our longstanding contacts on the logic of short-sight games: Section 4 of this paper owes a lot to his many suggestions and observations. This work is supported by the China Scholarship Council and NSFC grant No.61472369.

References

- [1] Alexandru Baltag, Sonja Smets & Jonathan A. Zvesper (2009): *Keep ‘hoping’ for rationality: a solution to the backward induction paradox*. *Synthese* 169(2), pp. 301–333, doi:10.1007/s11229-009-9559-z.
- [2] Johan van Benthem (2011): *Exploring a Theory of Play*. In: *Proc. of TARK*, pp. 12–16, doi:10.1145/2000378.2000380.
- [3] Johan van Benthem (2014): *Logic in Games*. MIT Press, doi:10.1080/00048402.2015.1063071.
- [4] Johan van Benthem & Amélie Gheerbrant (2010): *Game Solution, Epistemic Dynamics and Fixed-Point Logics*. *Fundam. Inform.* 100(1-4), pp. 19–41.
- [5] Johan van Benthem, Sieuwert Van Otterloo & Olivier Roy (2006): *Preference logic, conditionals, and solution concepts in games*. In: H. Lagerlund, S. Lindström, and R. Sliwinski, editors, *Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg*, University of Uppsala, pp. 61–76.
- [6] Johan van Benthem, Eric Pacuit & Olivier Roy (2011): *Toward a Theory of Play: A Logical Perspective on Games and Interaction*. *Games* 2(1), pp. 52–86, doi:10.3390/g2010052.
- [7] J. L. Bermúdez (2009): *Decision Theory and Rationality*. Oxford University Press, doi:10.1093/acprof:oso/9780199548026.001.0001.
- [8] Patrick Blackburn, Maarten de Rijke & Yde Venema (2001): *Modal logic*. Cambridge University Press, doi:10.1017/CBO9781107050884.
- [9] Blai Tirant Bonet Bretto (2004): *Modeling and Solving Sequential Decision Problems with Uncertainty and Partial Information*. Ph.D. thesis, Department of Computer Science, University of California, Los Angeles.
- [10] Cedric Degremont, Soumya Paul & Nicholas Asher (2014): *A Logic of Sights*. *Journal of Logic and Computation*.
- [11] Hans van Ditmarsch, Wiebe van der Hoek & Barteld Kooi (2007): *Dynamic Epistemic Logic*. *Synthese library* 337, Springer, doi:10.1007/978-1-4020-5839-4.
- [12] D. Edwards & T. Hart (1963): *The Alpha-Beta Heuristic*. Technical Report 30, MIT.
- [13] Yossi Feinberg (2012): *Games with Unawareness*. *Stanford Graduate School of Business Paper No. 2122*.
- [14] Davide Grossi & Paolo Turrini (2012): *Short sight in extensive games*. In: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pp. 805–812. Available at <http://dl.acm.org/citation.cfm?id=2343776.2343812>.
- [15] Joseph Y. Halpern & Leandro C. Rêgo (2009): *Reasoning About Knowledge of Unawareness Revisited*. In: *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge, TARK '09*, ACM, New York, NY, USA, pp. 166–173, doi:10.1145/1562814.1562838.
- [16] Joseph Y. Halpern & Leandro C. Rêgo (2014): *Extensive Games with Possibly Unaware Players*. *Mathematical Social Sciences* 70, pp. 42–58, doi:10.1016/j.mathsocsci.2012.11.002.
- [17] Joseph Y. Halpern & Leandro Chaves Rêgo (2006): *Extensive games with possibly unaware players*. In: *AAMAS*, pp. 744–751, doi:10.1145/1160633.1160765.
- [18] Sven Ove Hansson (1994): *Decision Theory -A Brief Introduction*.
- [19] Paul Harrenstein, Wiebe van der Hoek, John-Jules Meyer & Cees Witteveen (2003): *A Modal Characterization of Nash Equilibrium*. *Fundam. Inf.* 57(2-4), pp. 281–321. Available at <http://dl.acm.org/citation.cfm?id=986485.986494>.
- [20] Paul Harrenstein, Wiebe Van Der Hoek, John jules Meyer & Cees Witteveen (2002): *On modal logic interpretations of games*. In: *Procs ECAI 2002*, pp. 28–32.
- [21] Aviad Heifetz, Martin Meier & Burkhard C. Schipper (2013): *Dynamic unawareness and rationalizable behavior*. *Games and Economic Behavior* 81, pp. 50–68, doi:10.1016/j.geb.2013.04.003.
- [22] W. van der Hoek & M. Pauly (2006): *Modal logic for games and information*. In: *Handbook of Modal Logic*, Elsevier, pp. 1078–1148, doi:10.1016/S1570-2464(07)80023-1.

- [23] Brett Houlding (2008): *Sequential Decision Making with Adaptive Utility*. Ph.D. thesis, Department of Mathematical Sciences, Durham University.
- [24] Kjetil Høyland & Stein W. Wallace (2001): *Generating Scenario Trees for Multistage Decision Problems*. *Management Science* 47(2), pp. pp. 295–307, doi:10.1287/mnsc.47.2.295.9834.
- [25] Yew Jin Lim & Wee Sun Lee (2006): *Properties of forward pruning in game-tree search*. In: *proceedings of the 21st national conference on Artificial intelligence - Volume 2, AAAI'06*, AAAI Press, pp. 1020–1025.
- [26] Michael Lederman Littman (1996): *Algorithms for Sequential Decision-making*. Ph.D. thesis, Brown University, Providence, RI, USA.
- [27] Chanjuan Liu, Fenrong Liu & Kaile Su (2013): *A Logic for Extensive Games with Short Sight*. In: *LORI*, pp. 332–336, doi:10.1007/978-3-642-40948-6_28.
- [28] D. Warner North (1968): *A tutorial introduction to decision theory*. *IEEE Transactions on Systems Science and Cybernetics*, doi:10.1109/TSSC.1968.300114.
- [29] Martin J Osborne (2004): *An Introduction to Game Theory*. 2, Oxford University Press.
- [30] Martin J Osborne & Ariel Rubinstein (1994): *A Course in Game Theory*. MIT Press. Available at <https://mitpress.mit.edu/books/course-game-theory>.
- [31] Andrés Perea (2014): *FROM CLASSICAL TO EPISTEMIC GAME THEORY*. *International Game Theory Review (IGTR)* 16(01), doi:10.1142/S0219198914400015.
- [32] Martin Peterson (2009): *An Introduction to Decision Theory*, 1 edition. *Cambridge Introductions to Philosophy*, Cambridge University Press, doi:10.1017/CBO9780511800917.
- [33] Francesca Rossi, Kristen Brent Venable & Toby Walsh (2011): *A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice*. Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool Publishers, doi:10.2200/S00372ED1V01Y201107AIM014.