Epistemic Game Theoretical Reasoning in History Based Models

Can BAŞKENT Guy McCUSKER

Department of Computer Science, University of Bath, England

can@canbaskent.net www.canbaskent.net/logic

Third International Workshop on Strategic Reasoning, Oxford
September 21-22, 2015

Outlook of the Talk

- History Based Models
- Adding Preferences
- Prisoners' Dilemma
- An Agree-to-Disagree type Result

Motivation

- Need for a formal structure that can discuss the **epistemic**, **temporal**, **preferential**, and **strategic** aspects of logico-game theoretical reasoning.

- Talk about the knowledge of the players, their turns and preferences (and their rationality) and how they strategize based on or not the **game history**.

- Discuss **security games** and **irrational/inconsistent preferences and strategies** of the players within this framework.

Background

History based structures, proposed by Parikh and Ramanujam (Parikh & Ramanujam, 2003), suggest a formal framework that lies between process models, interpreted systems and propositional dynamic logics; and suggest replacing the state-based epistemic models with histories which are seen as sequences of events.

They have been used to model epistemic messages and communication between agents, deontic obligations and the relation between obligations and knowledge (Parikh & Ramanujam, 2003; Pacuit, 2007; Pacuit *et al.*, 2006). Moreover, they are similar to interpreted systems (Fagin *et al.*, 1995; Pacuit, 2007).

Basic Logical Structure I

History based models are constructed by using a given set of events **E** and agents **A**. Events can be seen as actions or moves which vary over time and affect the knowledge of the agents.

A finite set of events is denoted as $\mathbf{E}^*$, and for each agent i, $\mathbf{E}_i \subseteq \mathbf{E}$ is the set of events which are "seen" by i. A finite sequence of events from **E** is denoted by lowercase h, whereas a possibly infinite sequence of events is denoted by uppercase H. We call them both *histories*.

Basic Logical Structure II

We denote the concatenation of finite history h with H by hH. For a set of events $\mathbf{E}$, $\mathcal{H}^*_{\mathbf{E}}$ denotes the set of all finite histories with events from $\mathbf{E}$ and $\mathcal{H}_{\mathbf{E}}$ denotes the set of all histories.

Given two histories H, H', $H \leq H'$ denotes that H is a prefix of H'. We denote the length of finite h with $\text{len}(h)$. For a history H, H_t denotes that $H_t \leq H$ with $\text{len}(H_t) = t$.

We define *global history* as a sequence of events, finite or infinite, where a *local history* is the history of a particular agent. For any set of histories $\mathcal{H}$, the set $\text{FinPre}(\mathcal{H})$ denotes the set of finite prefixes of the histories in $\mathcal{H}$.

Basic Logical Structure III

Given an agent i and a global history H, the agent i can only access some of H. For two histories H, H', if the agent can access to the same parts of H and H', then H and H' are indistinguishable for i.

A function $\lambda_i : \text{FinPre}(H) \to \mathbf{E}_i^*$ is called a *locality function* for agent i and a global history H. Based on locality functions, the epistemic indistinguishability $\sim_i$ for agent i is defined between two histories H, H' as follows: If $H \sim_i H'$, then $\lambda_i(H) = \lambda_i(H')$.

Basic Logical Structure IV

The locality function as given above is rather general. First, we assume that agents' clock is consistent with the global clock. Second, $\lambda_i(H)$ is embeddable in H, that is the events in $\lambda_i(H)$ appear in H in the same order. In other words "agents are not wrong on about the events that they witness" (Pacuit, 2007).

For obvious reasons, $\sim_i$ is an equivalence relation. Thus, the epistemic logic of history based structures is the standard multi-agent epistemic logic $\mathbf{S5}_n$.

Syntax and Semantics I

Given a set **P** of propositional letters, the syntax of history based models can be given as follows

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid \bigcirc\varphi \mid \varphi U \varphi$$

where $p \in \mathbf{P}$ and $i \in \mathbf{A}$. The epistemic modality for agent i is K_i and the operator $\bigcirc$ is the next-time modality. We call U the *until operator*.

A history based model is given as a tuple $M = \{\mathcal{H}, \mathbf{E}_1, \ldots, \mathbf{E}_n, \lambda_1, \ldots, \lambda_n, \mathcal{V}\}$ where $\mathcal{V}$ is a valuation function which is defined in the standard fashion as follows: $V : \text{FinPre}(\mathcal{H}) \mapsto \wp(\mathbf{P})$.

Syntax and Semantics II

History based models semantically evaluates formulas at history-time pairs. At history H and time t in a model M, the satisfaction of a formula φ is denoted as $H, t \models_M \varphi$, and defined inductively as follows.

$$
\begin{array}{lll}
H, t \models_M p & \text{iff} & H_t \in V(p), \\
H, t \models_M \neg\varphi & \text{iff} & H, t \not\models_M \varphi, \\
H, t \models_M \varphi \wedge \psi & \text{iff} & H, t \models_M \varphi \text{ and } H, t \models_M \psi, \\
H, t \models_M \bigcirc\varphi & \text{iff} & H, t + 1 \models_M \varphi, \\
H, t \models_M K_i\varphi & \text{iff} & \forall H' \in \mathcal{H} \text{ and } H_t {\sim_i} H'_t \text{ implies } H', t \models_M \varphi, \\
H, t \models_M \varphi U \psi & \text{iff} & \exists k \geq t \text{ such that } H, k \models_M \psi \text{ and} \\
 & & \forall l,\ t \leq l < k \text{ implies } H, l \models_M \varphi.
\end{array}
$$

Syntax and Semantics III

The axioms for history based models are given as follows.

- All tautologies of propositional logic,
- $K_i(\varphi \rightarrow \psi) \rightarrow (K_i\varphi \rightarrow K_i\psi)$,
- $K_i\varphi \rightarrow \varphi \wedge K_iK_i\varphi$,
- $\neg K_i\varphi \rightarrow K_i\neg K_i\varphi$,
- $\bigcirc(\varphi \rightarrow \psi) \rightarrow (\bigcirc\varphi \rightarrow \bigcirc\psi)$,
- $\bigcirc\neg\varphi \leftrightarrow \neg \bigcirc \varphi$,
- $\varphi U\psi \leftrightarrow \psi \vee (\varphi \wedge \bigcirc(\varphi U\psi))$.

The rules of inference are modus ponens, and normalization for all three modalities.

Syntax and Semantics IV

Bisimulations can be defined as follows in a rather tedious way.

For history based models M, M', a bisimulation $\bowtie$ between M and M' is a tuple $\bowtie = (\bowtie_0, \bowtie_1)$ where $\bowtie_0 \subseteq M \times M'$ and $\bowtie_1 \subseteq M^2 \times M'^2$ such that

Propositional base case:

- If $H, t \bowtie_0 H', t'$, then H, t and H', t' satisfy the same propositional variable,

Temporal forth case:

- If $H, t \bowtie_0 H', t'$ and $t < u$, then there is u' in M' such that $t' < u'$, $H, u \bowtie_0 H', u'$ and $(H, t), (H, u) \bowtie_1 (H', t'), (H', u')$,

Syntax and Semantics V

- If $(H,t),(H,u) \bowtie_1 (H',t'),(H',u')$ and if there is v' with $t' < v' < u'$, then there exists v such that $t < v < u$ and $H,v \bowtie_0 H',v'$,

Temporal back case:

- If $H,t \bowtie_0 H',t'$ and $t' < u'$, then there is u in M such that $t < u$, $H,u \bowtie_0 H',u'$ and $(H,t),(H,u) \bowtie_1 (H',t'),(H',u')$,
- If $(H,t),(H,u) \bowtie_1 (H',t'),(H',u')$ and if there is v with $t < v < u$, then there exists v' such that $t' < v' < u'$ and $H,v \bowtie_0 H',v'$,

Epistemic forth case:

- If $H,t \bowtie_0 H',t'$ and $H_t \sim_i K_l$, then there is K',l' in M' such that $K,l \bowtie_0 K',l'$ and $H'_{t'} \sim_i K'_{l'}$,

Syntax and Semantics VI

Epistemic back case:

- If $H, t \bowtie_0 H', t'$ and $H'_{t'} \sim_i K_{l'}$, then there is K, l in M such that $K, l \bowtie_0 K', l'$ and $H_t \sim_i K_l$,

Theorem

For history based models M, M', if $M \bowtie M'$, then they satisfy the same formula.

Preferences over Histories I

For an agent i, and histories H, H', the expression $H \preceq_i H'$ denotes that "the agent i (weakly) prefers H' to H".

The preference relation is taken as a pre-order satisfying reflexivity and transitivity.

We can amend the syntax of the logic of history based models with the modal operator $\Diamond_i \varphi$ which expresses that there is a history which is at least as good as the current one and satisfies φ for agent i, with the following formal semantics.

$$H, t \models \Diamond_i \varphi \text{ iff } \exists H'.H \preceq_i H' \text{ and } H', t \models \varphi$$

We call the logic of history based structures with preference modality as HBPL after *history based preference logic*.

Preferences over Histories II

We take $\Diamond_i$ as an S4 modality with the usual rules of inference.

HBPL can be supplemented with various additional axioms to express some other interactive epistemic, temporal and game theoretical properties. Here we consider a few.

Connectedness of Preferences The connectedness property for the preference relation suggests that any two histories are comparable. Therefore, it can be formalized as $\forall H, H'.H \preceq_i H' \lor H' \preceq_i H$. The modal axiom that corresponds to it is the following axiom: $\Box_i(\Box_i\varphi \to \psi) \lor \Box_i(\Box_i\psi \to \varphi)$. This renders the frame with preference modality as a total pre-order.

Preferences over Histories III

Epistemic Perfect Recall The agents with perfect recall retain knowledge once they acquired it. The standard axiom for this property is given as follows: $K_i \bigcirc \varphi \to \bigcirc K_i \varphi$ which is valid in HBPL.

Preferential Perfectness By *preferential perfectness*, we mean that agents do not change their preferences in time. Consider the scheme $\Box_i \bigcirc \varphi \to \bigcirc \Box_i \varphi$ which is also valid in HBPL.

Prisoners' Dilemma I

Viewed as histories with imposed subjectives preferences, HBPL is helpful in formalizing epistemic games. As an application, we consider how HBPL computes best responses in Prisoners' Dilemma (PD, for short).

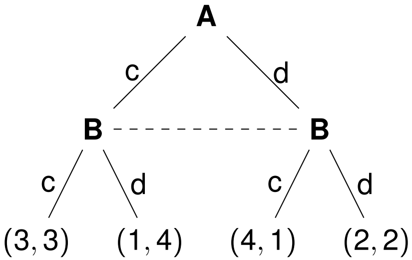

Figure: *Extensive form representation of PI*

Prisoners' Dilemma II

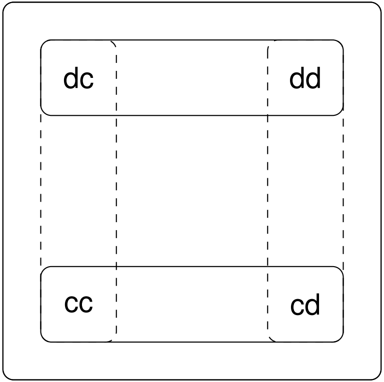

Figure: *Equivalence classes of histories for PI*

Analysis of Prisoners' Dilemma I

Let us consider PD in its extensive normal form where the utility pair (u_A, u_B) denotes the utility of the players **A** and **B** respectively.

Epistemic indistinguishability of the states for player **B** is denoted by the dashed line. In the history xy, the first event denotes Player **A**'s move while the second one denotes Player **B**'s move.

Due to the utilities associated with the players, we have cc $\preceq_B$ cd and dc $\preceq_B$ dd. Similarly, cc $\preceq_A$ dc and cd $\preceq_A$ dd.

We define best response of agent i in a two-player game as follows where $-i$ denotes the players other than i.

$$BR_i = \sim_{-i} \cap \preceq_i$$

Analysis of Prisoners' Dilemma II

Now let us see how we can verify the best responses of the players.

Recall that for both players, the best response is defect (the move d).

Analysis of Prisoners' Dilemma III

We start with Player **A.**

cc, $t \not\models BR_A$	since there is dc such that dc $\sim_B$ cc and cc $\preceq_A$ dc,
cd, $t \not\models BR_A$	since there is dd such that dd $\sim_B$ cd and cd $\preceq_A$ dd,
dc, $t \models BR_A$	as the only alternative cc fails to bring a higher utility,
dd, $t \models BR_A$	as the only alternative cd fails to bring a higher utility.

Analysis of Prisoners' Dilemma IV

Similarly for player **B**:

cc, $t \not\models BR_B$	since there is cd such that cd $\sim_A$ cc and cc $\preceq_A$ cd,
dc, $t \not\models BR_B$	since there is dd such that dd $\sim_A$ dc and dc $\preceq_A$ dd,
cd, $t \models BR_B$	as the only alternative cc fails to bring a higher utility,
dd, $t \models BR_B$	as the only alternative dc fails to bring a higher utility.

Analysis of Prisoners' Dilemma V

Based on the above analysis, Nash equilibrium can be observed at dd which is the state where neither of the agents can unilaterally benefit by diverging from. If **A** diverges, then the history cd is obtained which is not preferable for him. Similarly, if **B** diverges, then the history dc is obtained which is not preferable for him either. Thus, dd is the Nash equilibrium of PD.

Agree-to-Disagree Result in HBPL - I

Let us remember briefly Aumann's celebrated result (Aumann, 1976).

The Annals of Statistics
1976, Vol. 4, No. 6, 1236–1239

AGREEING TO DISAGREE[1]

BY ROBERT J. AUMANN

Stanford University and the Hebrew University of Jerusalem

Two people, 1 and 2, are said to have *common knowledge* of an event E if both know it, 1 knows that 2 knows it, 2 knows that 1 knows is, 1 knows that 2 knows that 1 knows it, and so on.

THEOREM. *If two people have the same priors, and their posteriors for an event A are common knowledge, then these posteriors are equal.*

If two people have the same priors, and their posteriors for a given event A are common knowledge, then these posteriors must be equal. This is so even though they may base their posteriors on quite different information. In brief, people with the same priors *cannot agree to disagree*.

Agree-to-Disagree Result in HBPL - II

In this section, we consider its reformulation by Dov Samet, and apply it to history based models.

Our application contains two different levels of complexity.

The first deals with the game play and constructs a history which includes the moves of all players and the local knowledge of players.

The second provides a global view of the model by forming equivalence classes of histories introducing additional meta-structure.

Agree-to-Disagree Result in HBPL - III

Let us now start with defining some standard epistemic operators following (Samet, 2010).

Definition

For a given set of agents **A** and a formula φ, we define $E_{\mathbf{A}}\varphi$ which reads "everyone in **A** knows φ". Formally, $E_{\mathbf{A}}\varphi = \bigwedge_{i \in \mathbf{A}} K_i\varphi$. We define the common knowledge operator $C_{\mathbf{A}}\varphi$ which reads "φ is common knowledge among **A**" as follows $C_{\mathbf{A}}\varphi = E_{\mathbf{A}}\varphi \wedge E^2_{\mathbf{A}}\varphi \wedge \cdots \wedge \ldots E^m_{\mathbf{A}}\varphi \wedge \ldots$, where $E^1_{\mathbf{A}} = E_{\mathbf{A}}\varphi$ and $E^{k+1}_{\mathbf{A}}\varphi = E_{\mathbf{A}}E^k_{\mathbf{A}}\varphi$, for $k \geq 1$.

Agree-to-Disagree Result in HBPL - IV

The epistemic indistinguishability relation $\sim_i$ for agent i makes it possible to redefine history based models as epistemic set models in a way that we can compare agents' knowledge relative to a given protocol (Samet, 2010). In order to achieve this, we define a set valued function which takes a set and returns a partition in that set that belongs to the agent. Given a protocol $\mathcal{H}$, we define $\kappa_i : 2^{\mathcal{H}} \mapsto 2^{\mathcal{H}}$. For simplicity, we will consider sets of finite histories, and denote the sets of histories with bold letters such as $\mathbf{h}, \mathbf{h}'$ etc. In this model, for each agent, there exists a partitioning of the given protocol $\mathcal{H}$.

Agree-to-Disagree Result in HBPL - V

Now, in a given model, let π_i denote the agent i's partitioning of the protocol $\mathcal{H}$. That is, for each i, there exists equivalence classes of histories in $\mathcal{H}$. Similarly, $\pi_i(h)$ denotes the partition for agent i that contains h. In other words, for an agent at history h, the histories in $\pi_i(h)$ are *indistinguishable*.

We define $\kappa_i(\mathbf{h}) = \{h : \pi_i(h) \subseteq \mathbf{h}\}$. Simply put, for a set of histories $\mathbf{h}$, the set $\kappa_i(\mathbf{h})$ includes all the histories h whose partitions are contained in $\mathbf{h}$. The operator κ_i is a set valued operator which will express agent's knowledge. In order to achieve this, we stipulate that κ_i satisfies the following three properties, for given sets of histories $\mathbf{h}, \mathbf{h}'$ (Fagin *et al.*, 1995).

Agree-to-Disagree Result in HBPL - VI

1. $\kappa_i(\mathbf{h} \cap \mathbf{h}') = \kappa_i(\mathbf{h}) \cap \kappa_i(\mathbf{h}')$
2. $\kappa_i(\mathbf{h}) \subseteq \mathbf{h}$
3. $-\kappa_i(\mathbf{h}) = \kappa_i(-\kappa_i(\mathbf{h}))$

where $-$ denotes the set theoretical complement. The above three property makes κ_i an epistemic operator where the first condition corresponds to normality, the second one to veridicality and the last one to introspection in the traditional sense. Similarly, a common knowledge operator $\mathfrak{c}$ can be defined for sets of histories to express the common knowledge modality $C_{\mathbf{A}}$.

Extending the preference relation in HBPL, it is possible to compare agents' knowledge relative to each other, given a set of histories.

Agree-to-Disagree Result in HBPL - VII

Definition

Define the set of histories $[j > i]^{\mathcal{H}}$ in which agent j is at least as knowledgeable as agent i with respect to a given set of protocols $\mathcal{H}$ as follows.

$$[j > i]^{\mathcal{H}} := \bigcap_{\mathbf{h} \in 2^{\mathcal{H}}} -\kappa_i(\mathbf{h}) \cup \kappa_j(\mathbf{h})$$

By a slight abuse of notation, we will denote the proposition whose extension is the set $[j > i]^{\mathcal{H}}$ by the same notation.

Agree-to-Disagree Result in HBPL - VIII

The following lemma expresses that the finer the partitions, the more the knowledge.

Lemma

$h \in [j > i]^{\mathcal{H}}$ iff $\pi_j(h) \subseteq \pi_i(h)$.

Next shows how the ordering of agents' knowledge and epistemic partitions relate to each other.

Lemma

$h \in \kappa_i([j > i])$ iff $\pi_i(h) = \bigcup_{h' \in \pi_i(h)} \pi_j(h')$.

Agree-to-Disagree Result in HBPL - IX

The vector $\delta = (\delta_1, \ldots, \delta_n)$ is called a decision profile for n agents. In this context, we consider D as *any* set of decisions, not necessarily probabilistic or propositional.

For a decision $d \in D$, we define the proposition $[\delta_i = d]^{\mathcal{H}}$ with its extension.

$$[\delta_i = d]^{\mathcal{H}} = \{H \in \mathcal{H} : \delta_i(H) = d \text{ for all } H \in \mathcal{H}\}$$

Similarly, we will use $[\delta_i = d]^{\mathcal{H}}$ to denote both the set and the proposition.

Agree-to-Disagree Result in HBPL - X

We assume each agent knows his decision (Samet, 2010):

$[\delta_i = d]^{\mathcal{H}} \subseteq \kappa_i([\delta_i = d]^{\mathcal{H}})$. In other words, agents agree with those agents who know better. Let us put it formally and more carefully as follows.

$$\kappa_i([j > i]^{\mathcal{H}} \cap [\delta_j = d]^{\mathcal{H}}) \subseteq [\delta_i = d]^{\mathcal{H}}$$

Above **Sure Thing Principle** suggests that if an agent j is at least as knowledgable as another agent i, and if j's decision is d, then i's decision is also d.

Agree-to-Disagree Result in HBPL - XI

An agent i is called *an epistemic dummy* if all the agents are at least as knowledgeable as i.

Definition A decision profile $\mathbf{d}$ in a model with a protocol $\mathcal{H}$ with n agents is expandable if for any additional epistemic dummy i, there exists a decision profile $\mathbf{d}'$ which satisfies the sure thing principle.

Agree-to-Disagree Result in HBPL - XII

Note that for an expandable decision profile **d** and dummy agent i, **d** and **d**′ agree on the decisions of agents who are not dummies. Expandable decision profiles play an important role for the following theorem, which we adopt from (Samet, 2010).

Theorem

If δ is an expandable decision profile in a model with a protocol $\mathcal{H}$ with n agents, then for any decisions $d_1, \ldots, d_n$ in D which are not identical, $C(\bigwedge_{i \leq n}[\delta_i = d_i]^{\mathcal{H}})$ is nowhere satisfiable, in other words $\mathfrak{c}(\bigcap_{i \leq n}[\delta_i = d_i]^{\mathcal{H}}) = \emptyset$.

Possible Further Applications

The theorem above provides some good handles for systems security policies. In systems' security, it can be seen that attacker's and defender's decisions cannot be the same for a successful attack or defence. Also, it is not enough that they will have different decisions, those decisions cannot be commonly known among them which is perhaps not-surprising for systems security.

Thanks for your attention!

Talk slides and the paper are available at

`www.CanBaskent.net/Logic`

References I

AUMANN, ROBERT J. 1976.

Agreeing to Disagree.

The Annals of Statistics, **4**(6), 1236–9.

FAGIN, RONALD, HALPERN, JOSEPH Y., MOSES, YORAM, & VARDI, MOSHE Y. 1995.

Reasoning About Knowledge.

MIT Press.

PACUIT, ERIC. 2007.

Some Comments on History Based Structures.

Journal of Applied Logic, **5**(4), 613–24.

PACUIT, ERIC, PARIKH, ROHIT, & COGAN, EVA. 2006.

The Logic of Knowledge Based Obligation.

Synthese, **149**(2), 311–341.

References II

PARIKH, ROHIT, & RAMANUJAM, R. 2003.

A Knowledge Based Semantics of Messages.

Journal of Logic, Language and Information, **12**(4), 453 – 467.

SAMET, DOV. 2010.

Agreeing to Disagree: The Non-probabilistic Case.

Games and Economic Behavior, **69**(1), 169–174.