

An abstract approach to reasoning about games with mistaken and changing beliefs

Benedikt Löwe^{1,2,3}, Eric Pacuit^{4*}

¹ Institute for Logic, Language and Computation, Universiteit van Amsterdam,
Plantage Muidergracht 24, 1018 TV Amsterdam, The Netherlands;
`bloewe@science.uva.nl`

² Department Mathematik, Universität Hamburg, Bundesstrasse 55, 20146 Hamburg,
Germany

³ Mathematisches Institut, Rheinische Friedrich-Wilhelms-Universität Bonn,
Berlingstraße 1, 53115 Bonn, Germany

⁴ Department of Computer Science, Stanford University, 353 Serra Mall, Stanford CA
94305, United States of America; `epacuit@stanford.edu`

May 7, 2008,

1 Introduction

Formalizing and understanding reasoning processes in multi-agent situations with imperfect information is one of the foremost tasks of modern artificial intelligence research. After its early phase (“good old-fashioned AI” or “GO-FAI” using a phrase of John Haugeland), artificial intelligence did not pay much attention to logic as a means towards achieving this end; logic and artificial intelligence grew increasingly apart. In the past years, this development has reversed: logic overcame its traditional focus and started to study phenomena of interaction and interactive reasoning. This is best represented in the research area called *Logic and Games* in which logicians used dynamic techniques in order to study behaviour in game and communication situations and in what Parikh calls *Social Software* [Par02]. These new developments allowed logicians to provide theoretical insight in the general project of understanding reasoning processes in multi-agent situations. As representative examples from the vast literature, we would like to mention van Ditmarsch’s complete epistemic formalization of the game Cluedo[®] [vDi00] and the Broersen-Dignum-Dignum-Meyer deontic logic of deadlines [BroDigDigMey04].

These development have triggered a fruitful interaction between the *Logic and Games* community and game and decision theorists who always had an interest in analyzing what constitutes rational behaviour under uncertainty. One

* The first author would like to thank Kobe University for the hospitality during his visit in March 2008 where the paper was thoroughly revised, partially supported by Grant-in-Aid for Scientific Research (C) 19540127 of the *Japan Society for the Promotion of Science*.

particular interesting encounter between logic and game theory is the use of belief revision techniques in the sense of [Gär92] as a means of analysis of games.¹ The game-theoretic analysis of rationality and the study of belief revision have in common that they have a normative hue; they are concerned with questions of what constitutes rational behaviour and what would be quality measures for rationality.

In this paper, we take a decidedly different route. We do not believe (as some of the proponents of GOFAI might have) that logic is the sole answer to deep and intriguing questions about human behaviour, but we think that it might be a useful tool in simulating and understanding it to a certain degree and in specifically restricted areas of application. Therefore, we do not aim to resolve the question of what rational behaviour in games with mistaken and changing beliefs is. Rather, we develop a formal and abstract framework that allows us to reason about behaviour in games with mistaken and changing beliefs leaving aside normative questions concerning whether the agents are behaving “rationally”, we focus on what the agents *do* in a game. In this paper, we are not concerned the reasoning process of the (ideal) economic agent; rather, our intended application is *artificial* agents. For example, autonomous agents interacting with a human user or with each other as part of a computer game or in a virtual world. Arguably, when such agents interact, the underlying epistemic and rationality assumptions are much less important than the *actual* reasoning process used by the agents. Thus we assume the players are “programmed” to execute the backward induction algorithm when it is their turn to move. Our model highlights what happens when such agents interact in the presence of changing beliefs and preferences.

The underlying assumption is that at each stage of the game, the agents use Zermelo’s backward induction algorithm in order to decide which action to perform next.² We assume that each player has his or her own set of beliefs about the preferences of the other players. It is with respect to these subjective

¹ The following is a list of relevant papers combining epistemic logic and game theory: [Aum99], [Sta98], [BonBat99], [Bra07], [Mon+97], [dBr04]. More specifically, the papers [Sta98], [Boa04], and [Per∞] have pointed out the importance of belief revision in the context of reasoning about solution concepts in games.

² Backwards induction is the most intuitive method for solving extensive games. It was first employed by Zermelo [Zer13] to show that the game of chess is determined and has since then been extended to infinite games [GalSte53] and found many applications in game theory, mathematics and computer science. Despite its naturalness, the rationality of backward induction has been questioned, for instance in Rosenthal’s centipede games [Ros81], and in fact, experimental results show that human players do not follow the backward induction strategy in these games [McKPa92,NagTan98]. Stemming from an interesting debate between Ken Binmore [Bin96,Bin97] and Robert Aumann [Aum95] about the underlying decision-theoretic assumptions of the backward induction algorithm, much of the literature has focused on the epistemic conditions (*e.g.*, *common knowledge*) and rationality assumptions (*e.g.*, the players maximize subjective expected utility) that *guarantee* the players follow the backward induction solution. *Cf.* [Per07] for a survey of the relevant lit-

preferences that the players execute the backward induction algorithm assuming that the preferences and beliefs of the players change over time. Questions about rationality and the choice of belief revision functions are left open to be determined in tune with the concrete application by those who wish to apply our framework.

In § 2, we give an example story of mistaken beliefs that describes how we imagine our formal setting to be applied. Then, in § 3, we give the definitions for our formal system and how to use this setting to get a backward induction solution. In the following § 4, we then apply our semantics to the story related in § 2 and give an analysis of it.

Our § 5 contains a discussion of related work and future projects. We discuss the advantages of our approach over existing approaches and indicate how it can be connected to the existing literature.

2 A story of reasoning with false beliefs

The following is a fictitious story in the style of a TV drama. The reader can imagine that this is the outline of a script. The reasoning processes referred to in the story can be made visible to the audience by monologues (Walter talking to himself in his car) or by conversations with some *confidant* or *confidante*.

Sue and Jeff have known each other for years. They studied computer science together in the 1980s, and both started their own software companies in the 1990s. Sue is married to Walter, an artist, and Jeff is married to Mary. In the past years, Sue and Mary have become best friends.

However, unbeknownst to Sue, her husband Walter and Mary have an affair. Walter, being absolutely dependant on the money of his wife, has no intention of leaving her at all, and wants to avoid that she gets to know about this affair at all costs. He believes that the fact that Sue and Mary are best friends acts as a safeguard for his secret: Mary must be aware that she will lose Sue as a friend if Sue finds out, and Mary clearly doesn't want that. So, Walter convinced himself that Mary will never ask him to separate from Sue or –even worse– tell Sue about the affair.

Mary on the other hand is rather unhappy with Jeff, and really wants to leave him. She believes that her friendship with Sue is robust enough to survive the fact that she has an affair with Walter. In her dreams, she imagines a nice future with Walter. She is convinced that if she presses Walter enough, he will finally leave Sue for her. She can make up with Sue afterwards.

One morning, she gives Walter an ultimatum: he should make up his mind and choose between her and Sue. Walter is ultimately confused: he must have misjudged Mary. Stuck in the traffic jam on his way to an

erature and § 5.2 for some ideas how to reconcile our framework with this line of research.

appointment with a potential client, his mind raced: If he chooses Mary, then Sue would know about their affair, and Mary would lose her best friend. What was Mary thinking? The only rational explanation that he could come up with was that Mary wanted to be with him so badly that she would give up her friendship with Sue for it and expected him to say ‘yes’ to her ultimatum... Obviously, for financial reasons, Walter couldn’t leave Sue. But he needed to be careful here: if he said no to Mary, would she tell Sue? No, he reasoned, since then she would lose both Sue and him which is definitely worse than just losing him. So, he’d be safe. Smiling, he used his cell phone to call Mary and tell her that he would not leave Sue.

When she hung up the phone, Mary was fuming with anger. Apparently, Walter wanted to stay with Sue. “Well, if that’s what he wants, then I must have completely misjudged him. I should cut my losses, and at least be honest to my best friend and tell her,” she reasoned, and acted accordingly.

And Mary was right in her judgement of Sue. The two women discussed the matter, and when Walter returned from his appointment in the afternoon, his paintings were standing on the front lawn of their house and the lock of the front door had been changed. Walter gazed emptyly at his paintings searching for a logician to help him to figure out what had happened.

We should stress that human beings have no problems in analysing an episode like this – with ease, they can make judgements like “Walter is wrong about his judgement of Sue and Mary; after the ultimatum, there was no chance of staying together with Sue anymore, but he could have saved his relationship with Mary hadn’t he misjudged his wife”.

For computational situations, we would like to be able to do the same within some preferably simple formal system. Being able to formally access the reasoning structure of episodes like this is crucial for the analysis of games with mistaken beliefs.

3 Formalization

Let I be the finite set of players whom we denote with boldface capital letters. If $\vec{\mathbf{P}} = \langle \mathbf{P}_0, \dots, \mathbf{P}_n \rangle$ is a finite sequence of players symbols, we write $\vec{\mathbf{P}}\mathbf{P}$ for the extension of the sequence by another player symbol \mathbf{P} , i.e.,

$$\vec{\mathbf{P}}\mathbf{P} := \langle \mathbf{P}_0, \dots, \mathbf{P}_n, \mathbf{P} \rangle.$$

A tree T is a finite set of nodes together with an edge relation (in which any two nodes are connected by exactly one path). Let root_T denote the root of the tree and $\text{tn}(T)$ denote the set of terminal nodes of T . We write $t \in T$ if t is a node in the tree T . If $t \in T$, let $\text{succ}_T(t)$ denote the set of immediate T -successors of

T . The **depth** of the tree T is the number of elements of a longest path in T , and we denote it by $\text{dp}(T)$.

An **extensive game form** is a tuple $\langle I, T, \mu \rangle$ where I is a set of players, T is a tree and μ is a **moving function**. That is,

$$\mu : T \setminus \text{tn}(T) \rightarrow I,$$

where $\mu(t) = \mathbf{P}$ indicates that it is \mathbf{P} 's move at node t . We call total orders \succeq on $\text{tn}(T)$ **preferences**. We fix in advance a finite set \mathcal{P} of preferences to be considered in the formal analysis. A map $\succeq : I \rightarrow \mathcal{P}$ is called a **preference profile** or **description** and we write $\succeq_{\mathbf{P}}$ for $\succeq(\mathbf{P})$. If $\langle I, T, \mu \rangle$ is an extensive game form, and \succeq is a preference profile (if $t_1 \succeq_i t_2$, we say “player i prefers the node t_1 over node t_2 ”), then we call

$$\langle I, T, \mu, \succeq \rangle$$

an **extensive game**. This model of a game is completely standard and discussions can be found in any game theory text (for example, *cf.* [OsRub94]).

From now on, we fix an extensive game form $\mathfrak{G} = \langle I, T, \mu \rangle$ and a finite set of preferences \mathcal{P} . We consider functions

$$S : T \times I^{\leq \text{dp}(T)} \rightarrow \mathcal{P}^I,$$

i.e., functions that assign a description $S(t, \vec{\mathbf{P}})$ to a node t in the tree and a finite string $\vec{\mathbf{P}}$ of player symbols (such as **ABAC** or also \emptyset), and call them **states**.³

We now interpret the description $S(t, \emptyset)$ as the **true state of affairs** at position t . Note that the true state of affairs can change – for instance, if the game reaches a certain position, players may change their preferences based on what happened in the game so far.

If $S(t, \vec{\mathbf{P}})$ is one of the descriptions defined by the state S , we interpret $S(t, \vec{\mathbf{P}}\mathbf{P})$ as player \mathbf{P} 's belief about $S(t, \vec{\mathbf{P}})$. For instance, $S(t, \mathbf{A})$ describes player \mathbf{A} 's belief about the true state of affairs at t , and $S(t, \mathbf{AB})$ describes player \mathbf{B} 's belief about player \mathbf{A} 's beliefs about the true state.

Note that this way, a state can be easily relativized to the subjective beliefs of a particular player. If S is a state and \mathbf{P} is one of the players, we can define

$$S_{\mathbf{P}}(t, \vec{\mathbf{P}}) := S(t, \vec{\mathbf{P}}\mathbf{P}).$$

The result $S_{\mathbf{P}}$ is essentially a state.⁴ A convenient way of presenting a state S (since it presents the belief change in a particularly perspicuous way) is by

³ The restriction of the length of the sequences to a number $\leq \text{dp}(T)$ is not essential here. It is just to ensure that states are finite objects. Our algorithm will never need the values of states for longer strings, so we can safely cut the strings at this length.

⁴ The only difference is that it is only defined for sequences of length $\leq \text{dp}(T) - 1$. This is due to our technical restriction discussed in Footnote 3.

drawing the corresponding **description trees** $t \mapsto S(t, \vec{\mathbf{P}})$ (for non-terminal nodes t); this is the way we shall present the state in our example in §4.

Our notion of state is very abstract in that it does not take into account any commonsense properties of belief or rationality. For instance, our definition of state allows that $S(t, \emptyset)(\mathbf{A}) = \sqsubseteq$, but $S(t, \mathbf{A})(\mathbf{A}) \neq \sqsubseteq$, *i.e.*, player \mathbf{A} has preferences that he doesn't believe he has; or $S(t, \mathbf{A})(\mathbf{P}) = \sqsubseteq$, but $S(t, \mathbf{AA})(\mathbf{P}) \neq \sqsubseteq$, *i.e.*, a violation of positive introspection; or $S(t, \mathbf{A})(\mathbf{B}) = \sqsubseteq$, for some position t inconsistent with \mathbf{B} having the preference \sqsubseteq . All of these properties might be conceivably useful in some applications, and can easily be excluded by additional axioms if they are not. For instance, if S is a state, we could call a player \mathbf{P} **introspective** in S if for all sequences $\vec{\mathbf{P}}$, we have $S_{\mathbf{P}\vec{\mathbf{P}}} = S_{\mathbf{P}\mathbf{P}\vec{\mathbf{P}}}$, and we can call a state S **introspective** if all players are introspective in S . Another useful property is that of egocentricity. In many applications, the agents assume that their beliefs are common belief of all players. This can be formalized as follows: a player \mathbf{P} is called **egocentric** in S if for all sequences $\vec{\mathbf{P}}$, we have $S_{\mathbf{P}} = S_{\mathbf{P}\vec{\mathbf{P}}}$, and we call a state S **egocentric** if all players are egocentric in S .

If $S(t, \vec{\mathbf{P}}) \neq S_{\mathbf{P}}(t, \vec{\mathbf{P}})$, this is an instance of **mistaken belief**; if for $t \neq t' \in T$, we have $S_{\mathbf{P}}(t, \vec{\mathbf{P}}) \neq S_{\mathbf{P}}(t', \vec{\mathbf{P}})$, we call it an instance of **changing belief**. Even though we will not use it in §4, we can also handle changing preferences: if $S(t, \emptyset)(\mathbf{P}) \neq S(t', \emptyset)(\mathbf{P})$, then player \mathbf{P} **changed preferences** between t and t' . If for all $t, t' \in T$ and \mathbf{P} , we have $S(t, \vec{\mathbf{P}}) = S_{\mathbf{P}}(t', \vec{\mathbf{P}})$, we call S a **state of correct and unchanging beliefs** which is an important special case, as it reconstructs the usual Zermelo backward induction solution (cf. Proposition 1).

If $\mathfrak{G} = \langle I, T, \mu \rangle$ is an extensive game form, a **game model** is a tuple $\langle \mathfrak{G}, S \rangle$ where S is a state. Given a game model $\langle \mathfrak{G}, S \rangle$, we can now fully analyze the game and predict its outcome (assuming that the players follow the backward induction solution). In order to do this analysis, we shall construct labellings $\ell_{S_{\vec{\mathbf{P}}}} : T \rightarrow \mathbf{tn}(T)$ where $\ell_{S_{\vec{\mathbf{P}}}}$ is interpreted as the subjective belief relative to $\vec{\mathbf{P}}$ of the outcome of the game if it has reached the node t . For instance, $\ell_{S_{\mathbf{A}}}(t) = t^* \in \mathbf{tn}(T)$, then player \mathbf{A} believes that if the game reaches t , the eventual outcome is t^* .

If t is a terminal node, we just let $\ell_U := t$ for all states U . In order to calculate the label of a node t controlled by player \mathbf{P} , we need the \mathbf{P} -subjective labels of all of its successors. More precisely: If $t \in T$, $\mu(t) = \mathbf{P}$ and we fix a state U , then we can define ℓ_U as follows: find the U -true preference of player \mathbf{P} , *i.e.*, $\sqsubseteq = U(t, \emptyset)(\mathbf{P})$. Then consider the labels $\ell_{U_{\mathbf{P}}}(t')$ for all immediate successors t' of t and pick the \sqsubseteq -maximal of these, say, t^* . Then $\ell_U(t) := t^*$. Concisely, $\ell_U(t)$ is the $U(t, \emptyset)(\mu(t))$ -maximal element of the set $\{\ell_{U_{\mu(t)}}(t'); t' \in \mathbf{succ}(t)\}$.

After we have defined all subjective labellings, the true run can be read off recursively. Since our labels are the terminal nodes, for each t with $\mu(t) = \mathbf{P}$ and S , there is a unique $t' \in \mathbf{succ}(t)$ such that $\ell_{S_{\mathbf{P}}}(t') = \ell_S(t)$. Starting from the root, take at each step the unique successor determined by $\ell_S(t)$ until you reach a terminal node. Note that this procedure will not follow the subjective tree labellings: player \mathbf{P} might believe that the outcome of the game is t^* when the game reaches t' , but the next move is determined by player $\mathbf{P}^* = \mu(t')$ who

will use his or her own subjective labelling $\ell_{S_{\mathbf{P}^*}}$ in order to determine the next successor.

Proposition 1 *If S is a state of correct and unchanging beliefs, then our algorithm reproduces the usual backward induction solution.*

Proof. An easy inductive argument shows that if S is a state of correct and unchanging beliefs, then for all $\bar{\mathbf{P}}$, we have $\ell_S = \ell_{S_{\bar{\mathbf{P}}}}$. We also know that for any \mathbf{P} , t and $U = S_{\bar{\mathbf{P}}}$, $U(t, \emptyset)(\mathbf{P})$ is the true preference of player \mathbf{P} , say, $\succeq_{\mathbf{P}}$.

Thus, for any non-terminal node t , $\ell_U(t)$ is the $\succeq_{\mu(t)}$ -maximal element of the set $\{\ell_U(t'); t' \in \text{succ}(t)\}$. This is exactly the definition of the backward induction labelling. q.e.d.

Proposition 2 *If S is a state and \mathbf{P} is egocentric in S , then for all sequences $\bar{\mathbf{P}}$, we have that $\ell_{S_{\mathbf{P}}} = \ell_{S_{\bar{\mathbf{P}}}}$.*

Proof. We prove $\ell_{S_{\mathbf{P}}}(t) = \ell_{S_{\bar{\mathbf{P}}}}(t)$ by induction on the depth of the node t . For terminal nodes t , the claim is true by definition. If t is nonterminal, then $\ell_{S_{\mathbf{P}}}(t)$ is the $S_{\mathbf{P}}(t, \emptyset)(\mu(t))$ -maximal element of the set $\{\ell_{S_{\mathbf{P}\mu(t)}}(t'); t' \text{ is an immediate successor of } t\}$, and $\ell_{S_{\bar{\mathbf{P}}}}(t)$ is the $S_{\bar{\mathbf{P}}}(t, \emptyset)(\mu(t))$ -maximal element of the set $\{\ell_{S_{\bar{\mathbf{P}}\mu(t)}}(t'); t' \text{ is an immediate successor of } t\}$. But since \mathbf{P} is egocentric, we have $S_{\mathbf{P}}(t, \emptyset)(\mu(t)) = S_{\bar{\mathbf{P}}}(t, \emptyset)(\mu(t))$ and $\ell_{S_{\mathbf{P}\mu(t)}}(t') = \ell_{S_{\bar{\mathbf{P}}\mu(t)}}(t') = \ell_{S_{\mathbf{P}}}(t')$ by induction hypothesis, and thus the sets of labels of successors are identical. q.e.d.

4 Analysis of the story

In the following, we now transform the story discussed in §2 into a state that our formal system can deal with. Actually, we do not need to give a full state but can restrict ourselves to those values of the function S that are relevant for the algorithm given in §3. Note that this analysis cannot be done in an automated way: a human theorist is needed to transform the natural language story into its formal counterpart:

In our story, there are three relevant agents $\{\mathbf{W}, \mathbf{M}, \mathbf{S}\}$. By analyzing the decision points in the story line (*i.e.*, those moments in which one of the agents makes a decision about an action), we can see that there are five of those, corresponding to the non-terminal nodes in our game tree. We see easily that there are six conceivably possible outcomes in the story (the terminal nodes in the game tree). The game tree that results from this analysis of the story is given in Figure 1. In order to have an easily recognizable notation for the six possible outcomes, we use the notation \odot to indicate that a relationship is intact and \bullet to indicate that it is ended, and list the terminal nodes by the status of the relationships in the order **Sue-Walter**, **Mary-Walter** and **Mary-Sue**;

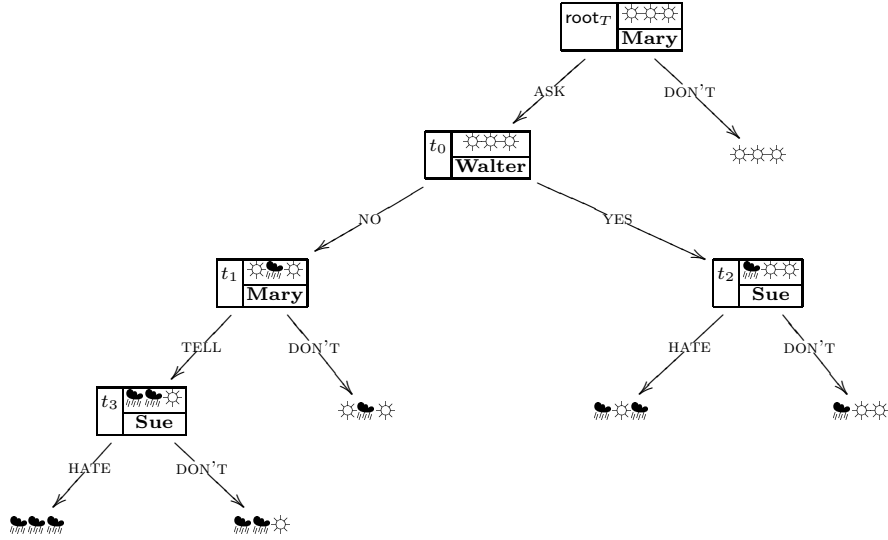



Fig. 1. The game tree for the story about Walter, Mary and Sue.

for instance, ☀☀☀☀☀ stands for “Walter left Sue, is together with Mary, and Sue hates Mary”. Since each outcome uniquely identifies one of the terminal nodes, we shall refer to the terminal nodes by these labels.

By analyzing the decision procedures in the story further, we realize that there are only five preferences relevant for the analysis of the story: the true preferences of all three players, and the mistaken beliefs of Walter about Sue’s preferences at the beginning of the game, and the mistaken beliefs of Mary about Walter’s preferences at the beginning of the game. We denote these preferences by \succ_W , \succ_M , \succ_S , \succ_{S^*} , and \succ_{W^*} , respectively:



In our situation, a description is a triple of preferences representing the preferences of Walter, Mary, and Sue. For instance, the true description at root_T is the triple

$$S(\text{root}_T, \emptyset) := \langle \succ_W, \succ_M, \succ_S \rangle.$$

As mentioned before, we shall give the state S mostly in terms of the description trees. We start with the true state of affairs in Figure 2. Note that in this tree, the description at each node is the same. This corresponds to the fact that there are no changing preferences in our example.

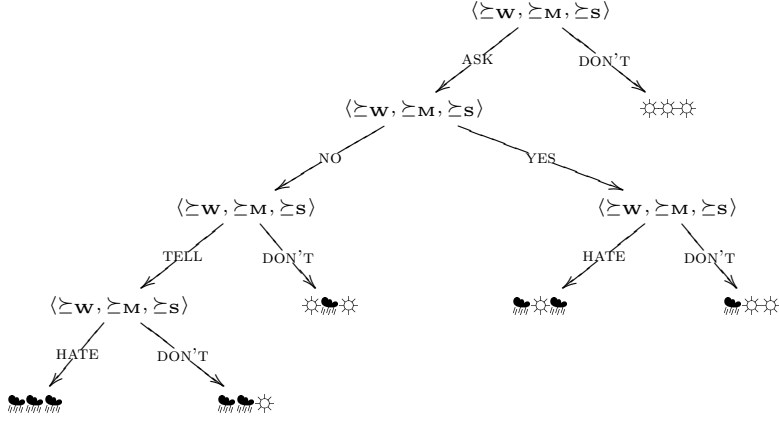


Fig. 2. The description tree $t \mapsto S(t, \emptyset)$ of true preferences

We continue with Walter’s beliefs. Walter’s first order belief is just $\langle \succeq_{\mathbf{w}}, \succeq_{\mathbf{M}}, \succeq_{\mathbf{s}^*} \rangle$ at each node: he is wrong about Sue, but correct about Mary. The belief revision that Walter undergoes in the story happens on the level of his second order beliefs: we give the description tree for $S_{\mathbf{WM}}$ in Figure 3 showing his belief revision at stage t_0 . He was originally egocentric, as $S_{\mathbf{WM}}(\text{root}_T, \emptyset) = S_{\mathbf{W}}(\text{root}_T, \emptyset)$, but realizes after Mary plays ASK that Mary cannot be right in her predictions of his own actions. He therefore corrects his beliefs about her beliefs to $\succeq_{\mathbf{w}^*}$ (which is what Mary actually believes). Unfortunately for him, he is still wrong about Mary’s belief about Sue.⁵ Note that we do not give a value for $S(t_3, \mathbf{WM})$: abstractly, we know that Walter must undergo yet another belief revision here, as Mary’s action TELL is inconsistent with the belief that Mary acts according to the preference $\succeq_{\mathbf{M}}$ and the belief that Sue according to $\succeq_{\mathbf{s}^*}$. However, in the story, we just learn that Walter is confused (“Walter gazed emptily at his paintings searching for a logician to help him to figure out what had happened”) and are thus not able to reconstruct a concrete belief revision for Walter at this step. Of course, since Walter does not move after t_3 , his beliefs are not relevant for the outcome of the game anymore.

The final description tree is that of Mary in Figure 4. She starts off with correct beliefs about Sue and false beliefs about Walter. After Walter plays NO in t_0 , Mary corrects her belief about Walter. For the calculation of the subjective labellings, we furthermore need the states $S_{\mathbf{MW}}$ and $S_{\mathbf{MWM}}$. The story suggests that Mary has no complicated iterated beliefs and that neither Walter nor Mary

⁵ One of the referees was confused by this point and wondered how Walter can be “right about Mary’s preferences” but wrong about his beliefs about what she’ll do. Of course, this is a crucial point in our algorithm. Knowing Mary’s preferences does not allow Walter to predict what she will do: since he is wrong about Sue’s preferences and (more importantly) about Mary’s beliefs about Sue’s preferences, he cannot correctly predict the outcome of the game in Mary’s subjective labelling.




Fig. 3. The description tree $t \mapsto S(t, \mathbf{WM}) = S_{\mathbf{WM}}(t, \emptyset)$ of Walter's beliefs about Mary's beliefs. The beliefs that cannot be extracted from the story are marked by “? ? ?”.

fail introspection. Therefore, we assume that Mary is egocentric and that Walter is introspective in the sense of § 3. Therefore, $S_{\mathbf{MW}} = S_{\mathbf{MWM}} = S_{\mathbf{M}}$.

Based on the description trees of S , $S_{\mathbf{M}}$, $S_{\mathbf{W}}$, $S_{\mathbf{MW}}$, $S_{\mathbf{WM}}$, and $S_{\mathbf{MWM}}$, we shall now proceed to compute the values of $\ell_U(t)$ for the relevant states U . Under an additional assumption of weakened egocentricity (see **Step 5**), we can calculate some additional values.

Step 1. Terminal nodes. By definition, for all states U , $\ell_U(t) = t$ if $t \in \mathbf{tn}(T)$.

Step 2. Nodes of depth one. There are two nodes of depth one, viz. t_2 and t_3 . In these cases, we can simply exploit the recursive definition and **Step 1** to get the following characterizations:

$$\ell_U(t_2) = \begin{cases} \bullet\bullet\bullet\star\star\star & \text{if } U(t_2, \emptyset)(\mathbf{S}) = \succeq_{\mathbf{S}} \\ \bullet\bullet\bullet\star\bullet\bullet & \text{if } U(t_2, \emptyset)(\mathbf{S}) = \succeq_{\mathbf{S}^*}, \end{cases} \quad (*)$$

$$\ell_U(t_3) = \begin{cases} \bullet\bullet\bullet\star & \text{if } U(t_3, \emptyset)(\mathbf{S}) = \succeq_{\mathbf{S}} \\ \bullet\bullet\bullet\bullet\bullet & \text{if } U(t_3, \emptyset)(\mathbf{S}) = \succeq_{\mathbf{S}^*}. \end{cases} \quad (**)$$

In particular, we get the following values for our six states under consideration:

$$\left. \begin{array}{l} \ell_S(t_2) \\ \ell_{S_{\mathbf{M}}}(t_2) \\ \ell_{S_{\mathbf{MW}}}(t_2) \\ \ell_{S_{\mathbf{MWM}}}(t_2) \end{array} \right\} = \bullet\bullet\bullet\star\star\star, \quad \left. \begin{array}{l} \ell_S(t_3) \\ \ell_{S_{\mathbf{M}}}(t_3) \\ \ell_{S_{\mathbf{MW}}}(t_3) \\ \ell_{S_{\mathbf{MWM}}}(t_3) \end{array} \right\} = \bullet\bullet\bullet\star,$$

$$\left. \begin{array}{l} \ell_{S_{\mathbf{W}}}(t_2) \\ \ell_{S_{\mathbf{WM}}}(t_2) \end{array} \right\} = \bullet\bullet\bullet\star\bullet\bullet, \quad \left. \begin{array}{l} \ell_{S_{\mathbf{W}}}(t_3) \\ \ell_{S_{\mathbf{WM}}}(t_3) \end{array} \right\} = \bullet\bullet\bullet\bullet\bullet.$$

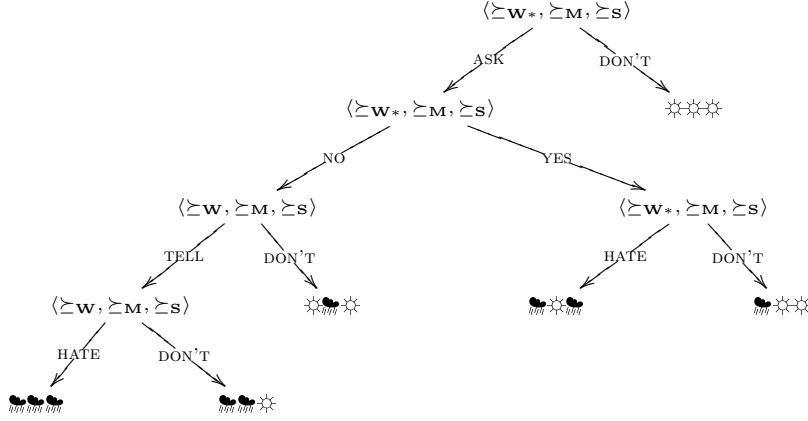


Fig. 4. The description tree $t \mapsto S(t, \mathbf{M}) = S_{\mathbf{M}}(t, \emptyset)$ of Mary's beliefs. This is also the description tree of $S_{\mathbf{MW}}$ and $S_{\mathbf{MWM}}$.

Step 3. Nodes of depth two. The node t_1 is the only node of depth two and it is a node controlled by Mary. Therefore, in order to calculate $\ell_U(t_1)$, we need the labels of the immediate successors of t_1 in the labelling $\ell_{U_{\mathbf{M}}}$. The immediate successors of t_1 are $\star\bullet\bullet\star$ and t_3 . By **Step 1**, $\ell_U(\star\bullet\bullet\star) = \star\bullet\bullet\star$ for all U . By **Step 2**, we know the value of $\ell_{U_{\mathbf{M}}}(t_3)$ for $U \in \{S, S_{\mathbf{W}}, S_{\mathbf{MW}}\}$. A glance at the table in **Step 2** gives us the values

$$\begin{aligned} \ell_{S_{\mathbf{M}}}(t_3) &= \ell_{S_{\mathbf{MWM}}}(t_3) = \bullet\bullet\bullet\star, \\ \text{and } \ell_{S_{\mathbf{WM}}}(t_3) &= \bullet\bullet\bullet\bullet. \end{aligned}$$

So, in the first two cases, we have to compare $\star\bullet\bullet\star$ and $\bullet\bullet\bullet\star$ according to the preference that the state U assigns to Mary. We see that $S(t_1, \emptyset)(\mathbf{M}) = S_{\mathbf{MW}}(t_1, \emptyset) = \succeq_{\mathbf{M}}$ and get

$$\ell_S(t_1) = \ell_{S_{\mathbf{MW}}}(t_1) = \bullet\bullet\bullet\star.$$

In the case of $S_{\mathbf{W}}$, we have to compare $\bullet\bullet\bullet\bullet$ and $\star\bullet\bullet\star$ according to $\succeq_{\mathbf{M}} = S_{\mathbf{W}}(t_1, \emptyset)(\mathbf{M})$, and get

$$\ell_{S_{\mathbf{W}}}(t_1) = \star\bullet\bullet\star.$$

Assuming Mary's introspection, we can use $\ell_{\mathbf{MM}}(t_3) = \bullet\bullet\bullet\star$ and calculate $\ell_{\mathbf{M}}(t_1) = \bullet\bullet\bullet\star$.

Step 4. Nodes of depth three. The only node of depth three is the node t_0 controlled by Walter. Its immediate successors are t_1 and t_2 , and thus we need $\ell_{U_{\mathbf{W}}}(t_1)$ and $\ell_{U_{\mathbf{W}}}(t_2)$ for a state U in order to calculate $\ell_U(t_0)$. Checking the calculations of **Step 2** (for t_2) and **Step 3** (for t_2), we see that we have determined these labels for $U \in \{S, S_{\mathbf{M}}\}$, and received the values

$$\begin{aligned} \ell_{S_{\mathbf{W}}}(t_1) &= \star\bullet\bullet\star, & \ell_{S_{\mathbf{MW}}}(t_1) &= \bullet\bullet\bullet\star, \\ \ell_{S_{\mathbf{W}}}(t_2) &= \bullet\bullet\star\star, & \ell_{S_{\mathbf{MW}}}(t_2) &= \bullet\bullet\star\star. \end{aligned}$$

For $U = S$, we compare $\star\star\star\star$ and $\star\star\star\star$ according to $\succeq_{\mathbf{W}}$ and get $\ell_S(t_0) = \star\star\star\star$, for $U = S_M$, we compare $\star\star\star\star$ and $\star\star\star\star$ according to $\succeq_{\mathbf{W}^*}$ and get $\ell_{S_M}(t_0) = \star\star\star\star$.

Using Walter's introspection (and thus using $\ell_{S_{\mathbf{W}\mathbf{W}}}(t_1) = \star\star\star\star$ and $\ell_{S_{\mathbf{W}\mathbf{W}}}(t_2) = \star\star\star\star$), we can also calculate $\ell_{S_{\mathbf{W}}}(t_0) = \star\star\star\star$.

Step 5. The root. The root is a node controlled by Mary, so we need to compare $\ell_{U_M}(t_0)$ and $\ell_{U_M}(\star\star\star\star) = \star\star\star\star$. We only have $\ell_{S_M}(t_0) = \star\star\star\star$, so we check that $\star\star\star\star \succeq_M \star\star\star\star$, and get $\ell_S(\text{root}_T) = \star\star\star\star$.

As before, under the assumption of Mary's introspection, we get that $\ell_{S_{\mathbf{M}\mathbf{M}}}(t_0) = \star\star\star\star$, and thus $\ell_{S_M}(\text{root}_T) = \star\star\star\star$.

To calculate $\ell_{S_{\mathbf{W}}}(\text{root}_T)$, we need to make the additional (but very mild and realistic) weakened egocentricity assumption of $S_{\mathbf{W}\mathbf{M}\mathbf{F}} = S_{\mathbf{W}\mathbf{M}}$. Then, going through **Step 1** to **Step 4** again, we can calculate $\ell_{S_{\mathbf{W}\mathbf{M}\mathbf{W}\mathbf{M}}}(t_3) = \star\star\star\star$, $\ell_{S_{\mathbf{W}\mathbf{M}\mathbf{W}}}(t_2) = \star\star\star\star$, $\ell_{S_{\mathbf{W}\mathbf{M}\mathbf{W}}}(t_1) = \star\star\star\star$, $\ell_{S_{\mathbf{W}\mathbf{M}}}(t_0) = \star\star\star\star$, and finally $\ell_{S_{\mathbf{W}}}(\text{root}_T) = \star\star\star\star$.

We finished our calculation of labels. It is useful to note again that the labelling functions behave quite differently from the usual backward induction labellings. In backward induction with common knowledge of the true preferences, you get one labelling (corresponding to our ℓ_S) that you can use to read off the moves of the game: we discussed this as a special case of our set-up in Proposition 1. In particular, each node has at least one successor with the same label. In our case, the labellings only indicate the belief about the eventual outcome of the player controlling the node. As a consequence, the labels of a node t correspond to player $\mu(t)$'s beliefs whereas the label of a successor node t' corresponds to $\mu(t')$'s belief.

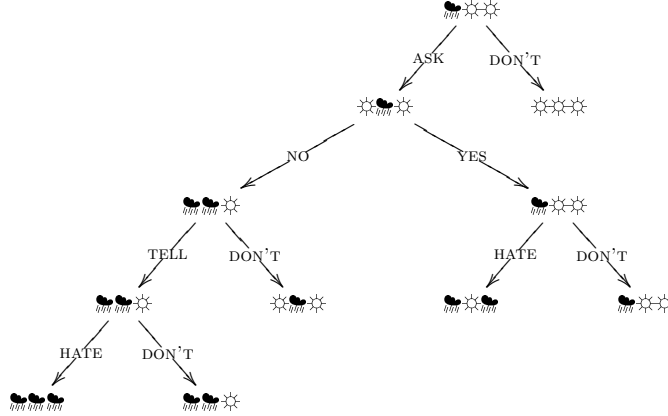


Fig. 5. The labelling $t \mapsto \ell_S(t)$.




Fig. 6. The (partial) labelling $t \mapsto \ell_{S_{\mathbf{W}}}(t)$. The label at t_0 has been computed under the assumption of introspection; the label at root_T has been computed with the additional assumption of $S_{\mathbf{WMP}} = S_{\mathbf{W}}\mathbf{M}$ (cf. **Step 5**).

Consider the labelling ℓ_S in Figure 5 where the nodes root_T and t_0 have no successor nodes with the same label. For instance, in order to understand $\ell_S(t_0)$, we have to look at the labelling $\ell_{S_{\mathbf{W}}}$ (as given in Figure 6), realizing that the value $\ell_S(t_0) = \star\bullet\star$ corresponds to the $\succeq_{\mathbf{W}}$ -maximal label among the $\ell_{S_{\mathbf{W}}}$ -values of the successors of t_0 .

We shall now use the labellings calculated in order to give an interpretation of the story and explain the course of events. We start in the node root_T which is a node in which Mary has to play, and consider Mary’s subjective labelling $\ell_{S_{\mathbf{M}}}$ as given in Figure 7. Mary believes that Walter will answer YES to her ultimatum, corresponding to “*She is convinced that if she presses Walter enough, he will finally leave Sue for her.*” In Figure 7, this is represented by $\ell_{S_{\mathbf{M}}}(t_0) = \bullet\star\star$. Mary’s belief that “*she can make up with Sue afterwards*”) is represented by $\ell_{S_{\mathbf{M}}}(t_3) = \bullet\bullet\star$ and $\ell_{S_{\mathbf{M}}}(t_2) = \bullet\star\star$. The label $\ell_S(\text{root}_T) = \bullet\star\star$ tells us that Mary plays ASK in her first move (“*One morning, she gives Walter an ultimatum: he should make up his mind and choose between her and Sue.*”).

In the story, we can now read the deliberations of Walter. He realizes that he had wrong beliefs about Mary (“*Walter is ultimately confused: he must have misjudged Mary*”) and revises his belief as discussed above. He now reasons about Mary’s actions if he plays NO as the next move: “*if he said no to Mary, would she tell Sue? No, he reasoned, since then she would lose both Sue and him which is definitely worse than just losing him.*” corresponds to $\ell_{S_{\mathbf{W}}}(t_1) = \star\bullet\star$. As a consequence, Walter plays NO in the second move (“*Smiling, he used his cell phone to call Mary and tell her that he would not leave Sue*”).

The third move is now Mary’s again who revises her beliefs about Walter (“*Well, if that’s what he wants, then I must have completely misjudged him.*”)




Fig. 7. The (partial) labelling $t \mapsto \ell_{S_M}(t)$. The value at t_1 has been computed under the assumption of introspection as in **Step 3**.

She follows her preferences and beliefs about Sue’s preferences: “*I should cut my losses, and at least be honest to my best friend and tell her,*” she reasoned, and acted accordingly.” Mary plays TELL in the third move.

Finally, Sue follows Mary’s predictions and plays DON’T: “*And Mary was right in her judgement of Sue. The two women discussed the matter, and when Walter returned from his appointment in the afternoon, his paintings were standing on the front lawn of their house and the lock of the front door had been changed.*”.

As part of our formal analysis, we can now also go beyond describing the behaviour of agents. For instance, the fact that Walter is surprised by Mary’s action ASK is represented by $\ell_{S_W}(\text{root}_T) = \ast\ast\ast\ast$ (“*Walter convinced himself that Mary will never ask him to separate from Sue*”). We can also do counterfactual reasoning about the storyline in terms of our labelling. This would become relevant for programming the reasoning of artificial agents based on our system (cf. §5.5).

5 Conclusion, related work and future Work

With the formal analysis of §§3 and 4, we fulfilled the goal mentioned at the end of §2: we have a formal system that allows to mimic the intuitive reasoning of human beings about the game situation. However, the definition of a formal system provides only the very first step. A lot of open questions and problems remain.

5.1 Related work

There are many other models for game-theoretic reasoning in the presence of uncertainty, and it would be wrong to claim that they cannot deal with the example given in this paper at least as well as our abstract framework. It should be stressed that the elaborate example in this paper is not more than that: an example to show how our algorithm works; we do not claim that our approach is the first one that can explain the behaviour of the agents in stories like this. We also do not claim that our algorithm adds to our understanding of game theory in general: game-theoretically, it is just the straightforward application of the basic idea of backward induction to the situation of changing and mistaken beliefs.

The benefit of our framework is its extraordinary simplicity: we make the player’s preferences the basic entities of the entire algorithm and encode the belief change into the notion of *state*, thus avoiding to have to discuss the belief change functions. Because of this, we get a very parsimonious and flexible algorithm that can be applied to many different situations. We disagree with one of the referees of this paper who says that “the choice of the type of belief is an intrinsic part of such a framework” — on the contrary, we believe that the flexibility of our approach is an asset.

Closest in spirit to our game-theoretic analysis is the excellent paper [Fei05] in which he provides a powerful language to talk about dynamic games. Clearly, Feinberg’s system is able to provide a clear analysis of our example story. Feinberg does provide notions of rationality for his system and is thus able to *derive* that the agents will follow the backward induction strategy under certain conditions [Fei05, § 4.2]. This relates closely to our § 5.2.

Our parsimonious language suggests an affinity with the dynamic logic solutions of logical puzzles. A number of modal logic approaches to reasoning about belief revision and preferences exist; among these, [vDiLab07], the *epistemic preference models* of [vBeLiu07] (based on Hansson’s [Han90] logics of preferences) and the proba-doxastic logic of [Auc05a,Auc05b] come closest to our approach. We discuss this relationship more closely in § 5.3.

5.2 Game-theoretic problems.

Our analysis presupposes backward induction. In order to connect our work to the large body of research on the foundations of rationality mentioned in § 1, one could turn this around by add new atomic formulae $[i, \Sigma]$ interpreted as “player i follows the strategy denoted by Σ ” and then analyze properties of the players (corresponding to rationality) that will ensure that players follow the subjective backward induction strategy. Adding more expressive power to our extremely frugal language will make our system resemble Feinberg’s *subjective framework* from [Fei05]. Under appropriate assumption, it is likely that analogues of Feinberg’s *subjective conditions on rationality implying backward induction* can be found for our system and analogues of Aumann’s classic theorem [Aum95] and Feinberg’s [Fei05, Proposition 10] can be proved.

5.3 The relationship to dynamic epistemic logic.

As indicated in § 5.1, the purely combinatorial character of our algorithm suggests a very close relation to the techniques of dynamic epistemic logic, and it is likely that there is a translation of our algorithm into the framework of dynamic epistemic logic.

Since backward induction can be expressed in dynamic epistemic logic [vBe01] and our algorithm is just based on a dynamic doxastic system and backward induction, it should be possible to do for our algorithm what van Ditmarsch did to belief revision in his [vDi05].

5.4 Analysis of story lines.

A conceptual system using our formal analysis could be used to analyze story lines. The question of what constitutes an interesting story has become more interesting for compute scientists as part of game design. Obviously, a large part of the entertainment that people get out of watching crime stories on TV or playing story-based games is unrelated to the actual plot of the stories, and rests in properties superficial to the actual story (such as the graphics animation in computer games or the quality of the actor in a TV drama). However, we believe that at least part of the entertainment derives from interesting epistemic story twists, either for the viewer or for the characters in the story.

Our simple formal analysis provides an abstract means of classifying storylines based on their epistemic and doxastic properties. An empirical project based on this would be to analyze a large sample of crime stories using our framework and find out whether certain formal patterns occur more frequently than others.

5.5 Game design.

The algorithm given in this paper could be used by artificial agents in computer games to mimic human behaviour. Epistemically more complicated computer games could feature artificial agents with more involved goals (such as trying to bring the human player to perform a certain action or trying to get a certain piece of information from the human player) that requires them to assess the doxastic situation.

An alternative use of our algorithm in game design could be connected to the project described above: if our formal model allows a classification of those epistemic story lines that human audiences find interesting as discussed in § 5.4, this information could be used to reverse engineer story lines (for instance for auto-generation of subplots in a computer game) in order to avoid story lines that are too straightforward or too complicated for the human players to enjoy.

References

- [Auc05a] Guillaume **Aucher**, A combined system for update logic and belief revision, *in*: Michael W. Barley, Nik Kasabov, Intelligent Agents

- and Multi-Agent Systems, 7th Pacific Rim Workshop on Multi-Agents, PRIMA 2004, Auckland, New Zealand, August 8–13, 2004, Revised Selected Papers, Springer-Verlag, 2005 [Lecture Notes in Artificial Intelligence 3371], pp. 1–17
- [Auc05b] Guillaume **Aucher**, How our beliefs contribute to interpret actions, *in*: Michal P. Petta, Laszlo Z. Varga (*eds.*), Multi-Agent Systems and Applications IV, 4th International Central and Eastern European Conference on Multi-Agent Systems, CEEMAS 2005, Budapest, Hungary, September 15–17, 2005, Proceedings, Springer-Verlag 2005 [Lecture Notes in Artificial Intelligence 3690], pp. 276–286
- [Aum95] Robert **Aumann**, Backward induction and common knowledge of rationality, **Games and Economic Behavior** 8 (1995), pp. 6–19
- [Aum99] Robert **Aumann**, Interactive epistemology I: knowledge, **International Journal of Game Theory** 28 (1999), pp. 263–300
- [BatSin99] Pierpaolo **Battigalli**, Marciano **Siniscalchi**, Hierarchies of conditional beliefs and interactive epistemology in dynamic games, **Journal of Economic Theory** 88 (1999), pp. 188–230
- [Bin96] Ken **Binmore**, A note on backward induction, **Games and Economic Behavior** 17 (1996), pp. 135–137
- [Bin97] Ken **Binmore**, Rationality and backward induction, **Journal of Economic Methodology** 4 (1997), pp. 23–41
- [Boa04] Oliver **Board**, Dynamic interactive epistemology, **Games and Economic Behavior** 49 (2004), pp. 49–80
- [BonBat99] Giacomo **Bonanno**, Pierpaolo **Battigalli**, Recent results on belief, knowledge and the epistemic foundations of game theory, **Research in Economics** 53 (1999), pp. 149–225
- [Bra07] Adam **Brandenburger**, The power of paradox: some recent developments in interactive epistemology, **International Journal of Game Theory** 35 (2007), pp. 465–492
- [BraKei01] Adam **Brandenburger**, H. Jerome **Keisler**, Epistemic conditions for iterated admissibility, *in*: Johan van Benthem (*ed.*), Proceedings of the 8th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2001), Certosa di Pontignano, University of Siena, Italy, July 8–10, 2001, Morgan Kaufman 2001, pp. 31–37
- [BroDigDigMey04] Jan **Broersen**, Frank **Dignum**, Virginia **Dignum**, John-Jules **Meyer**, Designing a Deontic Logic of Deadlines, *in*: Alessio Lomuscio, Donald Nute (*eds.*), Deontic Logic in Computer Science, 7th International Workshop on Deontic Logic in Computer Science, DEON 2004, Madeira, Portugal, May 26–28, 2004, Proceedings, Springer-Verlag, 2004, pp. 43–56
- [dBr04] Boudewijn **de Bruin**, Explaining Games: On the Logic of Game Theoretic Explanations, PhD Thesis, Universiteit van Amsterdam 2004, **ILLC Publications** DS-2004-03
- [Fag+95] Ronald **Fagin**, Joseph Y. **Halpern**, Yoram **Moses**, Moshe Y. **Vardi**, Reasoning about Knowledge, MIT Press, 1995
- [Fei05] Yossi **Feinberg**, Subjective reasoning — dynamic games, **Games and Economic Behavior** 52 (2005), pp. 54–93
- [GalSte53] David **Gale**, Frank M. **Stewart**, Infinite Games with Perfect Information, *in*: Harold W. Kuhn, Albert W. Tucker (*eds.*), Con-

- tributions to the Theory of Games II, Princeton 1953 [Annals of Mathematical Studies 28], pp. 245–266
- [Gär92] Peter **Gärdenfors**, Belief revision: an introduction, *in*: Peter Gärdenfors (*ed.*), *Belief Revision*, Cambridge University Press 1992 [Cambridge Tracts in Theoretical Computer Science 29], pp. 1–28
- [Han90] Sven Ove **Hansson**, Preference-based deontic logic (PDL), **Journal of Philosophical Logic** 19 (1990), pp. 75–93
- [McKPal92] Richard D. **McKelvey**, Thomas R. **Palfrey**, An experimental study of the centipede game, **Econometrica** 60 (1992), pp. 803–836
- [Mon+97] Philippe **Mongin**, Michael O. L. **Bacharach**, Louis-André **Gérard-Varet**, Hyun Song **Shin** (*eds.*), *Epistemic logic and the theory of games and decisions*, Kluwer Academic Publishers, 1997 [Theory and Decision Library, Series C: Game Theory, Mathematical Programming and Mathematical Economics 20]
- [NagTan98] Rosemarie **Nagel**, Fang Fang **Tang**, An experimental study on the centipede game in normal form — an investigation on learning, **Journal of Mathematical Psychology** 42 (1998), pp. 356–384
- [OsRub94] Martin J. **Osborne**, Ariel **Rubinstein**, *A Course in Game Theory*, MIT Press, 1994.
- [Par02] Rohit **Parikh**, Social Software, **Synthese** 132 (2002), pp. 187–211
- [Par03] Rohit **Parikh**, Levels of knowledge, games, and group action, **Research in Economics** 57 (2003), p. 267–281
- [Per07] Andrès **Perea**, Epistemic Foundations for Backward Induction: An Overview, *in*: Johan van Benthem, Dov Gabbay, Benedikt Löwe (*eds.*), *Interactive Logic, Selected Papers from the 7th Augustus de Morgan Workshop*, London, Amsterdam University Press, 2007 [Texts in Logic and Games 1], pp. 159–193
- [Per∞] Andrès **Perea**, A model of minimal probabilistic belief revision, *to appear in Theory and Decision*
- [Ros81] Robert W. **Rosenthal**, Games of perfect information, predatory pricing and the chain store, **Journal of Economic Theory** 25 (1981), pp. 92–100
- [Sta98] Robert **Stalnaker**, Belief revision in games: forward and backward induction, **Mathematical Social Sciences**, 36 (1998), pp. 31–56
- [vBe01] Johan **van Benthem**, Games in dynamic-epistemic logic, **Bulletin of Economic Research** 53 (2001), pp. 219–248
- [vBeLiu07] Johan F. A. K. **van Benthem**, Fenrong **Liu**, Dynamic Logic of Preference Upgrade, **Journal of Applied Non-Classical Logic** 17 (2007), pp. 157–182
- [vDi00] Hans Pieter **van Ditmarsch**, Knowledge games, PhD Thesis, Rijksuniversiteit Groningen, **ILLC Publications** DS-2000-06
- [vDi05] Hans **van Ditmarsch**, Belief change and dynamic logic, *in*: James Delgrande, Jérôme Lang, Hans Rott, Jean-Marc Tallon (*eds.*), *Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics*, Proceedings of Dagstuhl Seminar 05321, Schloß Dagstuhl, 2005
- [vDiLab07] Hans **van Ditmarsch**, Willem **Labuschagne**, My beliefs about your beliefs — a case study in theory of mind and epistemic logic, **Synthese** 155 (2007), pp. 191–209

[Zer13]

Ernst **Zermelo**, Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels, *in*: E. W. Hobson, A. E. H. Love (eds.), Proceedings of the Fifth International Congress of Mathematicians, Cambridge 1912, Volume 2, Cambridge 1913, pp. 501–504